

# Runtime Verification

Grigore Roşu

University of Illinois at Urbana-Champaign



# Contents

<b>1</b>	<b>Introduction</b>	<b>7</b>
1.1	A Taxonomy of Runtime Analysis Techniques . . . . .	18
1.1.1	Trace Storing versus Non-Storing Algorithms . . . . .	18
1.1.2	Synchronous versus Asynchronous Monitoring . . . . .	20
1.1.3	Predictive versus Exact Analysis . . . . .	21
<b>2</b>	<b>Background, Preliminaries, Notations</b>	<b>23</b>
2.1	Preliminaries . . . . .	26
2.1.1	Membership Equational Logic . . . . .	26
2.1.2	Maude . . . . .	30
<b>3</b>	<b>Safety Properties</b>	<b>37</b>
3.1	Finite Traces . . . . .	40
3.2	Infinite Traces . . . . .	47
3.3	Finite and Infinite Traces . . . . .	50
3.4	“Always Past” Characterization . . . . .	53
<b>4</b>	<b>Monitoring</b>	<b>57</b>
4.1	Specifying Safety Properties as Monitors . . . . .	57
4.2	Complexity of Monitoring a Safety Property . . . . .	62
4.3	Monitoring Safety Properties is Arbitrarily Hard . . . . .	68
4.4	Canonical Monitors . . . . .	70
<b>5</b>	<b>Event/Trace Observation</b>	<b>73</b>
<b>6</b>	<b>Monitor Synthesis: Finite State Machines (FSM)</b>	<b>75</b>
6.1	Binary Transition Trees (BTT) . . . . .	75
6.1.1	Multi-Transition and Binary Transition Tree Finite State Machines . . . . .	75

6.1.2	From MT-FSMs to BTT-FSMs . . . . .	80
<b>7</b>	<b>Monitor Synthesis: Extended Regular Expressions (ERE)</b>	<b>85</b>
7.1	Monitoring ERE Safety Needs Non-Elementary Space . . . . .	85
7.1.1	Discussion and Relevance of the Lower-Bound Result . . . . .	86
7.1.2	The Lower-Bound Result . . . . .	91
7.2	Generating Optimal Monitors for ERE . . . . .	98
7.2.1	Introduction . . . . .	99
7.2.2	Extended Regular Expressions and Derivatives . . . . .	102
7.2.3	Hidden Logic and Coinduction . . . . .	104
7.2.4	Behavioral Equivalence, Satisfaction and Specification . . . . .	105
7.2.5	Generating Minimal DFA Monitors by Coinduction . . . . .	111
7.2.6	Implementation and Evaluation . . . . .	113
<b>8</b>	<b>Monitor Synthesis: Linear Temporal Logic (LTL)</b>	<b>117</b>
8.1	Finite Trace Future Time Linear Temporal Logic . . . . .	117
8.1.1	Finite Trace Semantics . . . . .	118
8.1.2	Finite Trace Semantics in Maude . . . . .	119
8.2	A Backwards, Asynchronous, but Efficient Algorithm . . . . .	121
8.2.1	An Example . . . . .	121
8.2.2	Generating Dynamic Programming Algorithms . . . . .	123
8.3	A Forwards and Often Synchronous Algorithm . . . . .	126
8.3.1	An Event Consuming Algorithm . . . . .	127
8.3.2	Correctness and Completeness . . . . .	131
8.3.3	Further Optimization by Memoization . . . . .	134
8.4	Generating Forwards, Synchronous and Efficient Monitors . . . . .	135
8.4.1	From LTL Formulae to BTT-FSMs . . . . .	136
8.4.2	Examples . . . . .	142
8.5	Conclusions . . . . .	143
<b>9</b>	<b>Efficient Monitoring of <math>\omega</math>-Languages</b>	<b>145</b>
9.1	Introduction . . . . .	145
9.2	Preliminaries: Büchi Automata . . . . .	149
9.3	Multi-Transitions and Binary Transition Trees . . . . .	150
9.4	Binary Transition Tree Finite State Machines . . . . .	153
9.5	Generating a <i>BTT-FSM</i> from a Büchi automaton . . . . .	153
9.6	Monitor Generation and MOP . . . . .	158
9.6.1	Evaluation . . . . .	159
9.7	Conclusions . . . . .	160

<b>10 More Stuff</b>	<b>165</b>
10.1 Optimal Monitoring of “Always Past” Temporal Safety . . . .	165
10.1.1 The Monitor Synthesis Algorithm . . . . .	165
10.1.2 A Maude Implementation of the Monitor Synthesizer .	168
<b>11 Parametric Property Monitoring</b>	<b>177</b>
<b>12 Predictive Runtime Analysis</b>	<b>179</b>
<b>13 Static Analysis to Improve Runtime Verification</b>	<b>181</b>
<b>14 Semantics-Based Runtime Verification</b>	<b>183</b>
14.1 Defining a Formal Semantics . . . . .	183
14.2 Semantics-Based Symbolic Execution . . . . .	183
14.3 Program Verification as Exhaustive Runtime Verification . . .	183
<b>15 Conclusion and Future Work</b>	<b>185</b>
15.1 Safety Properties and Monitoring . . . . .	185

Topics to cover:

- Safety properties and their monitoring. How many safety properties are there? Can they all be monitored? Complexity of monitoring in different formalism: LTL, RE, ERE, CFG, SRS. Both dynamic properties, like above, and static properties (e.g., complex heap patterns).
- Event/Trace observation. How to observe the execution of a program? Instrumentation vs. runtime environment.
- Monitor synthesis. Generating optimal monitors for several formalisms: LTL, FT/PT-LTL, RE, ERE, CFG, SRS, PT-CaRet, Allen TL.
- Parametric property monitoring. How to deal with multiple instances of monitors?
- Predictive runtime analysis. Vector clock vs SMT-based techniques.
- Static analysis to improve runtime verification. Improve runtime overhead by not instrumenting what is unnecessary. Improve prediction capability by looking beyond the trace.
- Semantics-based Runtime Verification. Defining a formal language semantics. Using a semantics to do symbolic execution and runtime verify properties. Ultimate goal: verify programs by exhaustive runtime verification.

Papers which have been completely absorbed:

- \* [118]: Section 7.2; Chapter 2
- \* [109]: Chapter 2; Chapter 3; Chapter 4; Section 7.1;
- \* [105]: Section 6.1; Chapter 8;

# Chapter 1

## Introduction

---

*Begin of intro stuff for chapter on safety*

---

---

**From SACS: *Abstract sacs:*** *This paper addresses the problem of runtime verification from a foundational perspective, answering questions like “Is there a consensus among the various definitions of a safety property?” (Answer: Yes), “How many safety properties exist?” (Answer: As many as real numbers), “How difficult is the problem of monitoring a safety property?” (Answer: Arbitrarily complex), “Is there any formalism that can express all safety properties?” (Answer: No), etc. Various definitions of safety properties as sets of execution traces have been proposed in the literature, some over finite traces, others over infinite traces, yet others over both finite and infinite traces. By employing cardinality arguments and a novel notion of persistence, this paper first establishes the existence of bijective correspondences between the various notions of safety property. It then shows that safety properties can be characterized as “always past” properties. Finally, it proposes a general notion of monitor, which allows to show that safety properties correspond precisely to the monitorable properties, and then to establish that monitoring a safety property is arbitrarily hard.*

---

---

**From safety:** *Abstract safety:* Various definitions of safety properties as sets of execution traces have been introduced in the literature, some over finite traces, others over infinite traces, yet others over both finite and infinite traces. By employing cardinality arguments, this paper first shows that these notions of safety are ultimately equivalent, by showing each of them to have the cardinal of the continuum. It is then shown that all safety properties can be characterized as “always past” properties, and then that the problem of monitoring a safety property can be arbitrarily hard. Finally, two decidable specification formalisms for safety properties are discussed, namely extended regular expressions and past time LTL. It is shown that monitoring the former requires non-elementary space. An optimal monitor synthesis algorithm is given for the latter; the generated monitors run in space linear with the number of temporal operators and in time linear with the size of the formula.

---

A *safety property* is a behavioral property which, once violated, cannot be satisfied anymore. For example, a property “always  $x > 0$ ” is violated when  $x \leq 0$  is observed for the first time; this safety property remains violated even though eventually  $x > 0$  might hold. That means that one can identify each safety property with a set of “bad” finite execution traces, with the intuition that once one of those is reached the safety property is violated.

There are several apparently different ways to formalize safety. Perhaps the most immediate one is to complement the “bad traces” above and thus to define a safety property as a prefix-closed property over finite traces (containing the “good traces”) – by “property” in this paper we mean a set of finite or infinite traces. Inspired by Lamport [85], Alpern and Schneider [5] define safety properties over infinite traces as ones with the property that if an infinite trace is unacceptable then there must be some finite prefix of it which is already unacceptable, in the sense that there is no acceptable infinite completion of it. Is there any relationship between these two definitions of safety? We show rather indirectly that there is, by showing that their corresponding sets of safety properties have the cardinal  $c$  of the continuum (i.e., the cardinal of  $\mathbb{R}$ , the set of real numbers), so there exists some bijective mapping between the two. Unfortunately, the existence of such a bijection is

as little informative as the existence of a bijection between the real numbers and the irrational numbers. To capture the relationship between finite- and infinite-trace safety properties in a meaningful way, we introduce a subset of finite-trace safety properties, called *persistent*, and then construct an explicit bijection between that subset and the infinite-trace safety properties. Interestingly, over finite traces there are as many safety properties as unrestricted properties (finite-traces are enumerable and  $\mathcal{P}(\mathbb{N})$  is in bijection with  $\mathbb{R}$ ), while over infinite traces there are  $c$  safety properties versus  $2^c$  unrestricted properties (infinite traces are in bijection with  $\mathbb{R}$ ).

It is also common to define safety properties as properties over both finite and infinite traces, the intuition for the finite traces being that of unfinished computations. For example, Lamport [86] extends the notion of infinite-trace safety properties to properties over both finite and infinite traces, while Schneider et al. [113, 54] give an alternative definition of safety over finite and infinite traces, called “execution monitoring”. One immediate technical advantage of allowing both finite and infinite traces is that one can define prefix-closed properties. We indirectly show that prefix-closeness is not a sufficient condition to define safety properties when infinite traces are also allowed, by showing that there are  $2^c$  prefix-closed properties versus, as expected, “only”  $c$  safety properties.

Another common way to specify safety properties is as “always past” properties, that is, as properties containing only words whose finite prefixes satisfy a given property. If  $P$  is a property on finite prefixes, then we write  $\Box P$  for the “always  $P$ ” safety property containing the words with prefixes in  $P$ . We show that specifying safety properties as “always past” properties is fully justified by showing that, for each of the three types of traces (finite, infinite, and both), the “always past” properties are precisely the safety properties as defined above. It is common to specify  $P$  using some logical formalism, for example past time linear temporal logic (past LTL) [89]; for example, one can specify “ $a$  before  $b$ ” in past LTL as the formula  $b \rightarrow \Diamond a$ .

The problem of monitoring safety properties is also investigated in this paper. Since there are as many safety properties as real numbers, it is not unexpected that some of them can be very hard to monitor. We show that the problem of monitoring a safety property is arbitrarily hard, by showing that it reduces to deciding membership of natural

numbers to a set of natural numbers. In particular, we can associate a safety property to any degree in the arithmetic hierarchy as well as to any complexity class in the decidable universe, whose monitoring is as hard as that degree or complexity class.

---

**From SACS:** *This paper makes three novel contributions, two technical and another pedagogical. On the technical side, it first introduces the notion of a persistent safety property, which appears to be the right finite-trace correspondent of an infinite-trace safety property, and uses it to show the cardinal equivalence of the various notions of safety property encountered in the literature. Also on the technical side, it rigorously defines the problem of monitoring a safety property, and it shows that it can be arbitrarily hard. On the pedagogical side, this paper offers the first comprehensive study and uniform presentation of safety properties and of their monitoring.*

---

---

**From safety:** *In practice not all ( $c = |\mathbb{R}|$ ) safety properties are meaningful, but only those ( $\aleph_0 = |\mathbb{N}|$ ) which are specifiable using formal specification languages or logics of interest. We also investigate the problem of monitoring safety properties expressed using two common formalisms, namely regular expressions extended with complement, also called extended regular expressions (ERE), and LTL. It is known that both formalisms allow polynomial finite-trace membership checking algorithms [70, 105] if one has random access to the trace, but that both require exponential space if the trace can only be analyzed online [103, 80]. It is also known that LTL can indeed be monitored in exponential space [29] and so is claimed<sup>1</sup> for EREs in [103]. We show that the claim in [103] is, unfortunately, wrong, by showing that ERE monitoring requires non-elementary space. To do so, we propose for any  $n \in \mathbb{N}$  a safety property  $P_n$  whose monitoring requires space non-elementary in  $n$ , as well as an ERE of size  $O(n^3)$ . Since the known monitoring algorithms for LTL in its full generality are asymptotically optimal, what is left to do is to consider important fragments of LTL. We focus on the “always past” fragment and give a monitor synthesis algorithm that takes formulae  $\varphi$  and generate monitors for them that need  $O(k)$  total space and  $O(|\varphi|)$  time to process each event, where  $k$  is the number of past operators in  $\varphi$ . This improves over the best known algorithm that needs space  $O(|\varphi|)$  (and same time).*

---



---

*End of intro stuff for chapter on safety*

---

---

from *J.ASE 2005: Techniques for efficiently evaluating future time Linear Temporal Logic (abbreviated LTL) formulae on finite execution traces* are presented. While the standard models of LTL are infinite traces, finite traces appear naturally when testing and/or monitoring real applications that only run for limited time periods. A finite trace variant of LTL is formally defined, together with an immediate executable semantics which turns out to be quite inefficient if used directly, via rewriting, as a monitoring procedure. Then three algorithms are investigated. First, a simple synthesis algorithm for monitors based on dynamic programming is presented; despite the efficiency of the generated monitors, they unfortunately need to analyze the trace backwards, thus making them unusable in most practical situations. To circumvent this problem, two rewriting-based practical algorithms are further investigated, one using rewriting directly as a means for online monitoring, and the other using rewriting to generate automata-like monitors, called binary transition tree finite state machines (and abbreviated BTT-FSMs). Both rewriting algorithms are implemented in Maude, an executable specification language based on a very efficient implementation of term rewriting. The first rewriting algorithm essentially consists of a set of equations establishing an executable semantics of LTL, using a simple formula transforming approach. This algorithm is further improved to build automata on-the-fly via caching and reuse of rewrites (called memoization), resulting in a very efficient and small Maude program that can be used to monitor program executions. The second rewriting algorithm builds on the first one and synthesizes provably minimal BTT-FSMs from LTL formulae, which can then be used to analyze execution traces online without the need for a rewriting system. The presented work is part of an ambitious runtime verification and monitoring project at NASA Ames, called PATHEXPLORER, and demonstrates that rewriting can be a tractable and attractive means for experimenting and implementing logics for program monitoring.

---



---

next is from *J.ASE 2005*

---

Future time Linear Temporal Logic, abbreviated LTL, was introduced by

Pnueli in 1977 [100] (see also [88, 90]) for stating properties about reactive and concurrent systems. LTL provides temporal operators that refer to the future/remaining part of an execution trace relative to a current point of reference. The standard models of LTL are infinite execution traces, reflecting the behavior of such systems as ideally always being ready to respond to requests. Methods, such as model checking, have been developed for proving programs correct with respect to requirements specified as LTL formulae. Several systems are currently being developed that apply model checking to software systems written in Java, C and C++ [11, 31, 68, 27, 99, 45, 126, 62, 128]. However, for very large systems, there is little hope that one can actually prove correctness, and one must in those cases rely on debugging and testing. In the context of highly reliable and/or safety critical systems, one would actually want to *monitor* a program execution during operation and to determine whether it conforms to its specification. Any violation of the specification can then be used to guide the execution of the program into a safe state, either manually or automatically. In this paper we describe a collection of algorithms for monitoring program executions against LTL formulae. It is demonstrated how term rewriting, and in particular the Maude rewriting system [24], can be used to implement some of these algorithms very efficiently and conveniently.

The work presented in this paper has been started as part of, and stimulated by, the PATHEXPLORER project at NASA Ames, and in particular the Java PATHEXPLORER (JPAX) tool [63, 64] for monitoring Java programs. JPAX facilitates automated instrumentation of Java byte-code, currently using Compaq's JTREK which is not public anymore, but soon using BCEL [28]. The instrumented code emits relevant events to an observer during execution (see Figure 1.1). The observer can be running a Maude [24] process as a special case, so Maude's rewriting engine can be used to drive a temporal logic operational semantics with program execution events. The observer may run on a different computer, in which case the events are transmitted over a socket. The system is driven by a specification, stating what properties to be proved and what parts of the code to be instrumented. When the observer receives the events it dispatches these to a set of observer modules, each module performing a particular analysis that has been requested. In addition to checking temporal logic requirements, modules have also been programmed to perform error pattern analysis of multi-threaded programs, predicting deadlock and data race potentials.

Using temporal logic in testing is an idea of broad practical and theoretical

Figure 1.1: Overview of JPAX .

interest. One example is the commercial Temporal Rover and DBRover tools [32, 33], in which LTL properties are translated into code, which is then inserted at chosen positions in the program and executed whenever reached during program execution. The MaC tool [87, 76] is another example of a runtime monitoring tool. Here, Java byte-code is automatically instrumented to generate events of interest during the execution. Of special interest is the temporal logic used in MaC, which can be classified as a past time interval logic convenient for expressing monitoring properties in a succinct way. All the systems above try to discharge the program execution events as soon as possible, in order to minimize the space requirements. In contrast, a technique is proposed in [78] where the execution events are stored in an SQL database at runtime and then analyzed by means of queries after the program terminates. The PET tool, described in [53, 52, 51], uses a future time temporal logic formula to guide the execution of a program for debugging purposes. Java MultiPathExplorer [117] is a tool which checks a past time LTL safety formula against a partial order extracted online from an execution trace. POTA [114] is another partial order trace analyzer system. Java-MoP [18] is a generic logic monitoring tool encouraging “monitoring-oriented programming” as a paradigm merging specification and implementation. Complexity results for testing a finite trace against temporal formulae expressed in different temporal logics are investigated in [92]. Algorithms using alternating automata to monitor LTL properties are proposed in [37], and a specialized LTL collecting statistics along the execution trace is described in [38]. Various algorithms to generate testing automata from temporal logic formulae are discussed in [101, 98], and [44] presents a Büchi automata inspired algorithm adapted to finite trace LTL.

The major goal of this paper is to present rewriting-based algorithms for effectively and efficiently evaluating LTL formulae on finite execution traces *online*, that is, by processing each event as it arrives. An important contribution of this paper is to show how a rewriting system, such as Maude, makes it possible to experiment with monitoring logics very efficiently and elegantly, and furthermore can be used as a practical program monitoring engine. This approach allows one to formalize ideas in a framework close to standard mathematics. The presented algorithms are considered in the

context of JPAX, but they can be easily adapted and used within other monitoring frameworks. We claim that the techniques presented in this paper, even though applied to LTL, are in fact generic and can be easily applied to other logics for monitoring. For example, in [103, 115] we applied the same generic, “formula transforming”, techniques to obtain rewriting based algorithms for situations in which the logic for monitoring was replaced by extended regular expressions (regular expressions with complement).

A non-trivial application of the rewriting based techniques presented in this paper is X9, a test-case generation and monitoring environment for a software system that controls the planetary NASA rover K9. This collaborative effort is described in more detail in [9] and it will be presented in full detail elsewhere soon. The rover controller, programmed in 35,000 lines of C++, essentially executes plans, where a plan is a tree-like structure consisting of actions and sub-actions. The leaf actions control various hardware on the rover, such as for example the camera and the wheels. The execution of a plan must cause the actions to be executed in the right order and must satisfy various time constraints, also part of the plan. Actions can start and eventually either terminate successfully or fail. Plans can specify how failed sub-actions can propagate upwards.

Testing the rover controller consists of generating plans and then monitoring that the plan actions are executed in the right order and that failures are propagated correctly. X9 automatically generates plans from a “grammar” of the structure of plans, using the Java PathFinder model checker [128]. For each plan, a set of temporal formulae that an execution of the plan must satisfy is also generated. For example, a plan may state that an action  $a$  should be executed by first executing a sub-action  $a_1$  and then a sub-action  $a_2$ , and that the failure of any of the sub-actions should not propagate: action  $a$  should eventually succeed, regardless of whether  $a_1$  or  $a_2$  fails. The generated temporal formulae will state these requirements, such as for example  $\Box(\text{start}(a) \rightarrow \langle \rangle \text{succeed}(a))$  saying that “it is always the case ( $\Box$ ) that when action  $a$  starts, then eventually ( $\langle \rangle$ ) it terminates successfully”, and execution traces are monitored against them.

X9 is currently being turned into a mature system to be used by the developer. It is completely automated, generating a web-page containing all the warnings found. The top-level web-page identifies all the test-cases that have failed (by violating some of the temporal properties), each linked to a web-page containing specifics such as the plan, the execution trace, and the properties that are violated. X9 has itself been tested by seeding

errors into the rover controller code. The automated monitoring relieves the programmer from manually analyzing printed execution traces. Extending the logic with real-time, as is planned in future work, is crucial for this application since plans are heavily annotated with time constraints.

In Section 1.1, based on our experience, we give a rough classification of monitoring and runtime analysis algorithms by considering three important criteria. A first criterion is whether the execution trace of the monitored or analyzed program needs to be stored or not. Storing a trace might be very useful for specific types of analysis because one could have random access to events, but storing an execution trace is an expensive operation in practice, so sometimes trace-storing algorithms may not be desirable. A second criterion regards the synchronicity of the monitor, more precisely whether the monitor is able to react as soon as the specification or the requirement has been violated. Synchronicity may often trigger running a validity checker for the logic under consideration, which is typically a very expensive task. Finally, monitoring and analysis algorithms can also be classified as “predictive” versus “exact”, where the “exact” ones monitor the observed execution trace as a flat list of events, while the predictive algorithms try to guess potential erroneous behaviors of programs that can occur under different executions. All the algorithms in this paper are exact.

This paper requires a certain amount of mathematical notions and notations, which we introduce in Section 2.1 together with Maude [24], a high-performance system supporting both membership equational logic [94] and rewriting logic [93]. The current version of Maude can do more than 3 million rewritings per second on standard PCs, and its compiled version is intended to support more than 15 million rewritings per second<sup>2</sup>, so it can quite well be used as an implementation language.

Section 8.1 defines the finite trace variant of linear temporal logic that we use in the rest of the paper. We found, by carefully analyzing several practical examples, that the most appropriate assumption to make at the end of the trace is that it is stationary in the last state. Then we define the semantics of the temporal operators using their usual meaning in infinite trace LTL, where the finite trace is infinitely extended by repeating the last state. Another option would be to consider that all atomic predicates are false or true in the state following the last one, but this would be problematic when inter-dependent predicates are involved, such as “gate-up” and “gate-down”.

---

<sup>2</sup>Personal communication by José Meseguer.

In previous work we described a technique which synthesizes efficient dynamic programming algorithms for checking LTL formulae on finite execution traces [104]. Even though this algorithm is not dependent on rewriting (but it could be easily implemented in Maude by rewriting as we did with its dual variant for past time LTL [59, 18]), for the sake of completeness we present it in some detail in Section 8.2. This algorithm evaluates a formula bottom-up for each point in the trace, going backwards from the final state towards the initial state. Unfortunately, despite its linear complexity, this algorithm cannot be used online because it is both asynchronous and trace-storing. In [60, 61, 59] we dualize this technique and apply it to past time LTL, in which case the trace more naturally can be examined in a forwards direction synchronously.

Section 8.3 presents our first practical rewriting-based algorithm, which can directly monitor an LTL formula. This algorithm originates in [55, 104] and it was partially presented at the Automated Software Engineering conference [57]. The algorithm is expressed as a set of equations establishing an executable semantics of LTL using a simple formula transforming approach. The idea is to rewrite or transform an LTL monitoring requirement formula  $\varphi$  when an event  $e$  is received, to a formula  $\varphi\{e\}$ , which represents the new requirement that the monitored system should fulfill for the remaining part of the trace. This way, the LTL formula to monitor “evolves” into other LTL formulae by subsequent transformations. We show, however, that the size of the evolving formula is in the worst-case exponentially bounded by the size of the original LTL formula, and also that an exponential space explosion cannot be avoided in certain unfortunate cases. The efficiency of this rewriting algorithm can be improved by almost an order of magnitude by caching and reusing rewrites (a.k.a. “memoization”), which is supported by Maude. This algorithm is often synchronous, though there are situations in which it misses reporting a violation at the exact event when it occurs. The violation is, however, detected at a subsequent event. This algorithm can be relatively easily transformed into a synchronous one if one is willing to pay the price of running a validity checker, like the one presented in Subsection 8.4.1, after processing each event. The practical result of Section 8.3 is a very efficient and small Maude program that can be used to monitor program executions. The decision to use Maude has made it very easy to experiment with logics and algorithms in monitoring.

We finally present an alternative solution to monitoring LTL in Section 8.4, where a rewriting-based algorithm is used to *generate* an optimal special

observer from an LTL formula. By optimality is meant everything one may expect, such as minimal number of states, forwards traversal of execution traces, synchronicity, efficiency, but also less standard optimality features, such as transiting from one state to another with a minimum amount of computation. In order to effectively do this we introduce the notion of *binary transition tree* (BTT), as a generalization of binary decision diagrams (BDD) [14], whose purpose is to provide an *optimal order* in which state predicates need to be evaluated to decide the next state. The motivation for this is that in practical applications evaluating a state predicate is a time consuming task, such as for example to check whether a vector is sorted. The associated finite state machines are called *binary transition tree finite state machines* (BTT-FSM). BTT-FSMs can be used to analyze execution traces without the need for a rewriting system, and can hence be used by observers written in traditional programming languages. The BTT-FSM generator, which includes a validity checker, is also implemented in Maude and has about 200 lines of code in total.

## 1.1 A Taxonomy of Runtime Analysis Techniques

A *runtime analysis technique* is regarded in a broad sense in this section; it can be a method or a concrete algorithm that analyzes the execution trace of a running program and concludes a certain property about that program. Runtime analysis algorithms can be arbitrarily complex, depending upon the kind of properties to be monitored or analyzed. Based on our experience with current procedures implemented in JPaX, in this section we make an attempt to classify runtime analysis techniques. The three criteria below are intended to be neither exhaustive nor always applicable, but we found them quite useful in practice. They are not specific to any particular logic or approach, so we present them before we introduce our logic and algorithms. In fact, this taxonomy will allow us to appropriately discuss the benefits and drawbacks of our algorithms presented in the rest of the paper.

### 1.1.1 Trace Storing versus Non-Storing Algorithms

As events are received from the monitored system, a runtime analysis algorithm typically maintains a state which allows it to reason about the monitored execution trace. Ideally, the amount of information needed to be stored by the monitor in its state depends only upon the property to be

monitored and *not* upon the number of already processed events. This is desired because, due to the huge amount of events that can be generated during a monitoring session, one would want one's monitoring algorithms to work in linear time with the number of events processed.

There are, however, situations where it is not possible or practically feasible to use storage whose size is a function of only the monitoring requirement. One example is that of monitoring *extended regular expressions* (ERE), i.e., regular expressions enriched with a complement operator. As shown by the success of scripting languages like PERL or PYTHON, software developers tend to understand and like regular expressions and feel comfortable to describe patterns using those, so ERE is a good candidate formalism to specify monitoring requirements (we limit ourselves to only patterns described via temporal logics in this paper though).

It is however known that ERE to automata translation algorithms suffer from a non-elementary state explosion problem, because a complement operation requires nondeterministic-to-deterministic automata conversions, which yield exponential blowups in the number of states. Since complement operations can be nested, generating automata from EREs may often not be feasible in practice. Fortunately, there are algorithms which avoid this state explosion problem, at the expense of having to store the execution trace and then, at the end of the monitoring session, to analyze it by traversing it forwards and backwards many times. The interested reader is referred to [70] for a  $O(n^3m)$  dynamic programming algorithm ( $n$  is the length of the execution trace and  $m$  is the size of the ERE), and to [130, 81] for  $O(n^2m)$  non-trivial algorithms.

Based on these observations, we propose a first criterion to classify monitoring algorithms, namely on whether they *store or do not store the execution trace*. In the case of EREs, trace storing algorithms are polynomial in the size of the trace and linear in the ERE requirement, while the non-storing ones are linear in the size of the trace and highly exponential in the size of the requirement. In this paper we show that trace storing algorithms for linear temporal logic can be linear in both the trace and the requirement (see Section 8.2), while trace non-storing ones are linear in the size of the trace but simply exponential in the size of the requirement.

Trace non-storing algorithms are apparently preferred, but, however, their size can be so big that it could make their use unamenable in certain important situations. One should carefully analyze the trade-offs in order to make the best choice in a particular situation.

### 1.1.2 Synchronous versus Asynchronous Monitoring

There are many safety critical applications in which one would want to report a violation of a requirement as soon as possible, and to not allow the monitored program to take any further action once a requirement is violated. We call this desired functionality *synchronous monitoring*. Otherwise, if a violation can only be detected after the monitored program executes several more steps or after it is stopped and its entire execution trace is needed to perform the analysis, then we call it *asynchronous monitoring*.

The dynamic programming algorithm presented in Section 8.2 is *not* synchronous, because one can detect a violation only after the program is stopped and its execution trace is available for backwards traversal. The algorithm presented in Section 8.3 is also asynchronous in general because there are universally false formulae which are detected so only at the end of an execution trace or only after several other execution steps. Consider, for example, that one monitors the finite trace LTL formula  $\neg \langle \rangle (\Box A \vee \Box \neg A)$ , which is false because at the end of any execution trace  $A$  either holds or not, or the formula  $\Box \Box A \wedge \Box \Box \neg A$ , which is also false but will be detected so only after two more events. However, the rewriting algorithm in Section 8.3 is synchronous in many practical situations. The algorithm in Section 8.4 is always synchronous, though one should be aware of the fact that its size may become a problem on large formulae.

In order for an LTL monitor to be synchronous, it needs to implement a validity checker for finite trace LTL, such as the one in Subsection 8.4.1 (Figure 8.2), and call it on the current formula after each event is processed. Checking validity of a finite trace LTL formula is very expensive (we are not aware of any theoretical result stating its exact complexity, but we believe that it is PSPACE-complete, like for standard infinite trace LTL [121]). We are currently considering providing a fully synchronous LTL monitoring module within JPAX, at the expense of calling a validity checker after each event, and let the user of the system choose either synchronous or asynchronous monitoring.

There are, however, many practical LTL formulae for which violation can be detected synchronously by the formula transforming rewriting-based algorithm presented in Section 8.3. Consider for example the sample formula of this paper,  $\Box (\text{green} \rightarrow \neg \text{red} \cup \text{yellow})$ , which is violated if and only if a red event is observed after a green one. The monitoring requirement of our algorithm, which initially is the formula itself, will not be changed unless a green event is received, in which case it will change to

$(\text{!red} \cup \text{yellow}) \wedge [](\text{green} \rightarrow \text{!red} \cup \text{yellow})$ . A yellow event will turn it back into the initial formula, a green event will keep it unchanged, but a red event will turn it into **false**. If this is the case, then the monitor declares the formula violated and appropriate actions can be taken. Notice that the violation was detected *exactly* when it occurred. A very interesting, practical and challenging problem is to find criteria that say when a formula can be synchronously monitored without the use of a validity checker.

### 1.1.3 Predictive versus Exact Analysis

An increasingly important class of runtime analysis algorithms are concerned with *predicting* anomalies in programs from *successful* observed executions. One such algorithm can be easily obtained by slightly modifying the *wait-for-graph* algorithm, which is typically used to *detect* when a system is in a deadlock state, to make it predict deadlocks. One way to do this is to *not* remove synchronization objects from the wait-for-graph when threads/processes release them. Then even though a system is not deadlock, a warning can be reported to users if a cycle is found in the wait-for-graph, because that represents a *potential* of a deadlock.

Another algorithm falling into the same category is Eraser [111], a datarace prediction procedure. For each shared memory region, Eraser maintains a set of *active locks* which protect it, which is intersected with the set of locks held by any accessing thread. If the set of active locks ever becomes empty then a warning is issued to the user, with the meaning that a potential unprotected access can take place. Both the deadlock and the datarace predictive algorithms are very successful in practice because they scale well and find many of the errors they are designed for. We have also implemented improved versions of these algorithms in Java PathExplorer.

We are currently also investigating predictive analysis of safety properties expressed using past time temporal logic, and a prototype system called Java MultiPathExplorer is being implemented [117]. The main idea here is to *instrument* Java classes to emit events timestamped by vector clocks [36], thus enabling the observer to extract a *partial order* reflecting the causal dependency on the memory accesses of the multithreaded program. If any linearization of that inferred partial order leads to a violation of the safety property then a warning is generated to the user, with the meaning that there can be executions of the multithreaded program, including the current one, which violate the requirements.

In this paper we restrict ourselves to only *exact* analysis of execution

traces. That means that the events in the trace are supposed to have occurred exactly in the received order (this can be easily enforced by maintaining a logical clock, then timestamping each event with the current clock, and then delivering the messages in increasing order of timestamps), and that we only check whether that particular order violates the monitoring requirements or not. Techniques for predicting future time LTL violations will be investigated elsewhere soon.

Although the taxonomy discussed in this section is intended to only be applied to tools, the problem domain may also admit a similar taxonomy. While such a taxonomy seems to be hard to accomplish in general, it would certainly be very useful because it would allow one to choose the proper runtime analysis technique for a given system and set of properties. However, like this paper shows, it is often the case that one can choose among several types of runtime analysis techniques for a given problem domain.

## Chapter 2

# Background, Preliminaries, Notations

---

*Add some structure to this chapter*

---

---

*from RV03 and RTA03*

---

We let  $\mathbb{N}$  denote the set of natural numbers including 0 but excluding the infinity symbol  $\infty$  and let  $\mathbb{N}_\infty$  denote the set  $\mathbb{N} \cup \{\infty\}$ . We also let  $\mathbb{Q}$  denote the set of rational numbers and  $\mathbb{R}$  the set of real numbers; as for natural numbers, the “ $\infty$ ” subscript can also be added to  $\mathbb{Q}$  and  $\mathbb{R}$  for the corresponding extensions of these sets.  $\mathbb{Q}^+$  and  $\mathbb{R}^+$  denote the sets of strictly positive (0 not included) rational and real numbers, respectively.

We fix a set  $\Sigma$  of elements called *events* or *states*. We call words in  $\Sigma^*$  *finite traces* and those in  $\Sigma^\omega$  *infinite traces*. If  $u \in \Sigma^* \cup \Sigma^\omega$  then  $u_i$  is the  $i$ -th state or event that appears in  $u$ . We call *finite-trace properties* sets  $P \subseteq \Sigma^*$  of finite traces, *infinite-trace properties* sets  $P \subseteq \Sigma^\omega$  of infinite traces, and just *properties* sets  $P \subseteq \Sigma^* \cup \Sigma^\omega$  of finite or infinite traces. If the finite or infinite aspect of traces is understood from context, then we may call any of the types or properties above just *properties*. We may write  $P(w)$  for a property  $P$  and a (finite or infinite) trace  $w$  whenever  $w \in P$ . Traces and properties are more commonly called *words* and *languages*, respectively, in the literature; we prefer to call them traces and properties to better reflect the intuition that our target application is monitoring and system observance, not formal languages. We take, however, the liberty to also call them words and languages whenever that terminology seems more appropriate.

In some cases states can be simply identified with their names, or labels, and specifications of properties on traces may just refer to those labels. For example, the regular expression  $(s_1 \cdot s_2)^*$  specifies all those finite traces starting with state  $s_1$  and in which states  $s_1$  and  $s_2$  alternate. In other cases, one can think of states as sets of atomic predicates, that is, predicates that hold in those states: if  $s$  is a state and  $a$  is an atomic predicate, then we say that  $a(s)$  is true iff  $a$  “holds” in  $s$ ; thus, if all it matters with respect to states is which predicates hold and which do not hold in each state, then states can be faithfully identified with sets of predicates. We prefer to stay loose with respect to what “holds” means, because, depending on the context, it can mean anything. In conventional software situations, atomic predicates can be: boolean expressions over variables of the program, their satisfaction being decided by evaluating them in the current state of the program; or whether a function is being called or returned from; or whether a particular variable is being written to; or whether a particular lock is being held by a particular thread; and so on. In the presence of atomic predicates, specifications of properties on traces typically only refer to the atomic predicates. For example, the property “always  $a$  before  $b$ ”, that is, those traces containing no state in which  $b$  holds that is not preceded by some state in which  $a$  holds (for example,  $a$  can stand for “authentication” and  $b$  for “resource access”), can be expressed in LTL as the formula  $\Box(b \rightarrow \Diamond a)$ .

Let us recall some basic notions and notations from formal languages, temporarily using the consecrated terminology of “words” and “languages” instead of traces and properties. For an alphabet  $\Sigma$ , let  $\mathcal{L}_\Sigma$  be the set of languages over  $\Sigma$ , i.e., the powerset  $\mathcal{P}(\Sigma^*)$ . By abuse of language and notation, let  $\emptyset$  be the empty language  $\{\}$  and  $\epsilon$  the language containing only the empty word,  $\{\epsilon\}$ . If  $L_1, L_2 \in \mathcal{L}_\Sigma$  then  $L_1 \cdot L_2$  is the language  $\{\alpha_1\alpha_2 \mid \alpha_1 \in L_1 \text{ and } \alpha_2 \in L_2\}$ . Note that  $L \cdot \emptyset = \emptyset \cdot L = \emptyset$  and  $L \cdot \epsilon = \epsilon \cdot L = L$ . If  $L \in \mathcal{L}_\Sigma$  then  $L^*$  is  $\{\alpha_1\alpha_2 \cdots \alpha_n \mid n \geq 0 \text{ and } \alpha_1, \alpha_2, \dots, \alpha_n \in L\}$  and  $\neg L$  is  $\Sigma^* - L$ .

We next recall some notions related to cardinality. If  $A$  is any set, we let  $|A|$  denote the *cardinal* of  $A$ , which expresses the size of  $A$ . When  $A$  is finite,  $|A|$  is precisely the number of elements of  $A$  and we call it a *finite cardinal*. Infinite sets can have different cardinals, called *transfinite* or even *infinite*. For example, natural numbers  $\mathbb{N}$  have the cardinal  $\aleph_0$  (pronounced “aleph zero”) and real numbers  $\mathbb{R}$  have the cardinal  $c$ , also called the *cardinal of the continuum*. Two sets  $A$  and  $B$  are said to have the same cardinal, written  $|A| = |B|$ , iff there is some bijective mapping between the two. We

write  $|A| \leq |B|$  iff there is some injective mapping from  $A$  to  $B$ .

The famous *Cantor-Bernstein-Schroeder theorem* states that if  $|A| \leq |B|$  and  $|B| \leq |A|$  then  $|A| = |B|$ . In other words, to show that there is some bijection between sets  $A$  and  $B$ , it suffices to find an injection from  $A$  to  $B$  and an injection from  $B$  to  $A$ . The two injections need not be bijections. For example, the inclusion of the interval  $(0, 1)$  in  $\mathbb{R}^+$  is obviously an injection, so  $|(0, 1)| \leq |\mathbb{R}^+|$ . On the other hand, the function  $x \mapsto x/(2x + 1)$  from  $\mathbb{R}^+$  to  $(0, 1)$  (in fact its codomain is the interval  $(0, 1/2)$ ) is also injective, so  $|\mathbb{R}^+| \leq |(0, 1)|$ . Neither of the two injective functions is bijective, yet by the Cantor-Bernstein-Schroeder theorem there is some bijection between  $(0, 1)$  and  $\mathbb{R}^+$ , that is,  $|(0, 1)| = |\mathbb{R}^+|$ . We will use this theorem to relate the various types of safety properties; for example, we will show that there is an injective function from safety properties over finite traces to safety properties over infinite traces and another injective function in the opposite direction. Unfortunately, the Cantor-Bernstein-Schroeder theorem is existential: it only says that some bijection exists between the two sets, but it does not give us an explicit bijection. Since the visualization of a concrete bijection between different sets of safety properties can be very meaningful, we will avoid using the Cantor-Bernstein-Schroeder theorem when we can find an explicit bijection between two sets of safety properties.

If  $A$  is a set of cardinal  $\alpha$ , then  $2^\alpha$  is the cardinal of  $\mathcal{P}(A)$ , the power set of  $A$  (the set of subsets of  $A$ ). It is known that  $2^{\aleph_0} = c$ , that is, there are as many sets of natural numbers as real numbers. The famous, still unanswered *continuum hypothesis*, states that there is no set whose size is strictly between  $\aleph_0$  and  $c$ ; more generally, it states that, for any transfinite cardinal  $\alpha$ , there is no proper cardinal between  $\alpha$  and  $2^\alpha$ . If  $A$  and  $B$  are infinite sets, then  $|A| + |B|$  and  $|A| \cdot |B|$  are the cardinals of the sets  $A \cup B$  and  $A \times B$ , respectively. An important property of transfinite cardinals is that of *absorption* – the larger cardinal absorbs the smaller one: if  $\alpha$  and  $\beta$  are transfinite cardinals such that  $\alpha \leq \beta$ , then  $\alpha + \beta = \alpha \cdot \beta = \beta$ ; in particular,  $c \cdot 2^c = 2^c$ . Besides sets of natural numbers, there are several other important sets that have cardinal  $c$ : streams (i.e., infinite sequences) of Booleans, streams of reals, non-empty closed or open intervals of reals, as well as the sets of all open or closed sets of reals, respectively (Exercise 1).

For our purposes, if  $\Sigma$  is an enumerable set of states, then  $\Sigma^*$  is also enumerable, so it has cardinal  $\aleph_0$ . Also, if  $|\Sigma| \leq c$ , in particular if it is finite, then  $\Sigma^\omega$  has the cardinal  $c$ , because it is equivalent to streams of states. We can then immediately infer that the set of finite-trace properties over  $\Sigma$  has

cardinal  $2^{\aleph_0} = c$ , while the set of infinite-trace properties has cardinal  $2^c$ .

---

*end from RV03 and RTA03*

---

## Exercises

**Exercise 1** *Show that each of the following sets have cardinal  $c$ : streams (i.e., infinite sequences) of Booleans; streams of natural numbers; streams of real numbers; closed intervals of real numbers; open intervals of real numbers; closed sets of real numbers; open sets of real numbers.*

---

*from J.ASE'05*

---

## 2.1 Preliminaries

In this section we recall notions and notations that will be used in the paper, including membership equational logic, term rewriting, Maude notation, and (infinite trace) linear temporal logics.

### 2.1.1 Membership Equational Logic

Membership equational logic (MEL) extends many- and order-sorted equational logic by allowing memberships of terms to sorts in addition to the usual equational sentences. We only recall those MEL notions which are necessary for understanding this paper; the interested reader is referred to [94, 13] for a comprehensive exposition of MEL.

#### Basic Definition

A *many-kinded algebraic signature*  $(K, \Sigma)$  consists of a set  $K$  and a  $(K^* \times K)$ -indexed set  $\Sigma = \{\Sigma_{k_1 k_2 \dots k_n, k} \mid k_1, k_2, \dots, k_n, k \in K\}$  of operations, where an operation  $\sigma \in \Sigma_{k_1 k_2 \dots k_n, k}$  is written  $\sigma : k_1 k_2 \dots k_n \rightarrow k$ . A *membership signature*  $\Omega$  is a triple  $(K, \Sigma, \pi)$  where  $K$  is a set of *kinds*,  $\Sigma$  is a  $K$ -kinded algebraic signature, and  $\pi : S \rightarrow K$  is a function that assigns to each element in its domain, called a *sort*, a kind. Therefore, sorts are grouped according to kinds and operations are defined on kinds. For simplicity, we will call a “membership signature” just a “signature” whenever there is no confusion.

For a *many-kinded signature*  $(K, \Sigma)$ , a  $\Sigma$ -algebra  $A$  consists of a  $K$ -indexed set  $\{A_k \mid k \in K\}$  together with interpretations of operations  $\sigma : k_1 k_2 \dots k_n \rightarrow k$  into functions  $A_\sigma : A_{k_1} \times A_{k_2} \times \dots \times A_{k_n} \rightarrow A_k$ . For any given signature  $\Omega = (K, \Sigma, \pi)$ , an  $\Omega$ -membership algebra  $A$  is a  $\Sigma$ -algebra together with a set  $A_s \subseteq A_{\pi(s)}$  for each sort  $s \in S$ . A particular algebra, called *term algebra*, is of special interest. Given a  $K$ -kinded signature  $\Sigma$  and a  $K$ -indexed set of *variables*  $X$ , let  $T_\Sigma(X)$  be the algebra of  $\Sigma$ -terms over variables in  $X$  extending  $X$  iteratively as follows: if  $\sigma : k_1 k_2 \dots k_n \rightarrow k$  and  $t_1 \in T_{\Sigma, k_1}(X)$ ,  $t_2 \in T_{\Sigma, k_2}(X)$ , ...,  $t_n \in T_{\Sigma, k_n}(X)$ , then  $\sigma(t_1, t_2, \dots, t_n) \in T_{\Sigma, k}(X)$ .

Given a signature  $\Omega$  and a  $K$ -indexed set of variables  $X$ , an *atomic*  $(\Omega, X)$ -equation has the form  $t = t'$ , where  $t, t' \in T_{\Sigma, k}(X)$ , and an *atomic*  $(\Omega, X)$ -membership has the form  $t : s$ , where  $s$  is a sort and  $t \in T_{\Sigma, \pi(s)}(X)$ . An  $\Omega$ -sentence in MEL has the form  $(\forall X) a \text{ if } a_1 \wedge \dots \wedge a_n$ , where  $a, a_1, \dots, a_n$  are atomic  $(\Omega, X)$ -equations or  $(\Omega, X)$ -memberships, and  $\{a_1, \dots, a_n\}$  is a set (no duplications). If  $n = 0$ , then the  $\Omega$ -sentence is called *unconditional* and written  $(\forall X) a$ . Equations are called *rewriting rules* when they are used only from left to right, as it will happen in this paper.

Given an  $\Omega$ -algebra  $A$  and a  $K$ -kinded map  $\theta : X \rightarrow A$ , then  $A, \theta \models_\Omega t = t'$  iff  $\theta(t) = \theta(t')$ , and  $A, \theta \models_\Omega t : s$  iff  $\theta(t) \in A_s$ .  $A$  satisfies  $(\forall X) a \text{ if } a_1 \wedge \dots \wedge a_n$ , written  $A \models_\Omega (\forall X) a \text{ if } a_1 \wedge \dots \wedge a_n$ , iff for each  $\theta : X \rightarrow A$ , if  $A, \theta \models_\Omega a_1$  and ... and  $A, \theta \models_\Omega a_n$ , then  $A, \theta \models_\Omega a$ .

An  $\Omega$ -specification (or  $\Omega$ -theory)  $T = (\Omega, E)$  in MEL consists of a signature  $\Omega$  and a set  $E$  of  $\Omega$ -sentences. An  $\Omega$ -algebra  $A$  satisfies (or is a model of)  $T = (\Omega, E)$ , written  $A \models T$ , iff it satisfies each sentence in  $E$ .

## Inference Rules

MEL admits complete deduction (see [94], where the rule of congruence is stated in a somewhat different but equivalent way). In the congruence rule below,  $\sigma \in \Sigma_{k_1 \dots k_i, k}$ ,  $W$  is a set of variables  $w_1 : k_1, \dots, w_{i-1} : k_{i-1}, w_{i+1} : k_{i+1}, \dots, w_n : k_n$ , and  $\sigma(W, t)$  is a shorthand for the term

$\sigma(w_1, \dots, w_{i-1}, t, w_{i+1}, \dots, w_n)$ :

- (1) Reflexivity : 
$$\frac{}{E \vdash_{\Omega} (\forall X) t = t}$$
- (2) Symmetry : 
$$\frac{E \vdash_{\Omega} (\forall X) t = t'}{E \vdash_{\Omega} (\forall X) t' = t}$$
- (3) Transitivity : 
$$\frac{E \vdash_{\Omega} (\forall X) t = t', E \vdash_{\Omega} (\forall X) t' = t''}{E \vdash_{\Omega} (\forall X) t = t''}$$
- (4) Congruence : 
$$\frac{E \vdash_{\Omega} (\forall X) t = t'}{E \vdash_{\Omega} (\forall X, W) \sigma(W, t) = \sigma(W, t'), \text{ for each } \sigma \in \Sigma}$$
- (5) Membership : 
$$\frac{E \vdash_{\Omega} (\forall X) t = t', E \vdash_{\Omega} (\forall X) t : s}{E \vdash_{\Omega} (\forall X) t' : s}$$
- (6) Modus Ponens : 
$$\left\{ \begin{array}{l} \text{Given a sentence in } E \\ (\forall Y) t = t' \text{ if } t_1 = t'_1 \wedge \dots \wedge t_n = t'_n \wedge w_1 : s_1 \wedge \dots \wedge w_m : s_m \\ \text{(resp. } (\forall Y) t : s \text{ if } t_1 = t'_1 \wedge \dots \wedge t_n = t'_n \wedge w_1 : s_1 \wedge \dots \wedge w_m : s_m) \\ \text{and } \theta : Y \rightarrow T_{\Sigma}(X) \text{ s.t. for all } i \in \{1, \dots, n\} \text{ and } j \in \{1, \dots, m\} \\ E \vdash_{\Omega} (\forall X) \theta(t_i) = \theta(t'_i), E \vdash_{\Omega} (\forall X) \theta(w_j) : s_j \\ \hline E \vdash_{\Omega} (\forall X) \theta(t) = \theta(t') \quad (\text{resp. } E \vdash_{\Omega} (\forall X) \theta(t) : s) \end{array} \right.$$

The rules above can therefore prove any unconditional equation or membership that is true in all membership algebras satisfying  $E$ . In order to derive conditional statements, we will therefore consider the standard technique adapting the “deduction theorem” to equational logics, namely deriving the conclusion of the sentence after adding the condition as an axiom; in order for this procedure to be correct, the variables used in the conclusion need to be first transformed into constants. All variables can be transformed into constants, so we only consider the following simplified rules:

- (7) Theorem of Constants : 
$$\frac{E \vdash_{\Omega \cup X} (\forall \emptyset) a \text{ if } a_1 \wedge \dots \wedge a_n}{E \vdash_{\Omega} (\forall X) a \text{ if } a_1 \wedge \dots \wedge a_n}$$
- (8) Implication Elimination : 
$$\frac{E \cup \{a_1, \dots, a_n\} \vdash_{\Omega} (\forall \emptyset) a}{E \vdash_{\Omega} (\forall \emptyset) a \text{ if } a_1 \wedge \dots \wedge a_n}$$

**Theorem 1** (from [94]) *With the notation above,  $E \models_{\Omega} (\forall X) a$  if  $a_1 \wedge \dots \wedge a_n$  if and only if  $E \vdash_{\Omega} (\forall X) a$  if  $a_1 \wedge \dots \wedge a_n$ . Moreover, any statement can be proved by first applying rule (7), then (8), and then a series of rules (1) to (6).*

This theorem is used within the correctness proof of the monitoring algorithm in Section 8.3.

### Initial Semantics and Induction

MEL specifications are often intended to allow only a restricted class of models (or algebras). For example, a specification of natural numbers defined using the Peano equational axioms would have many “undesired” models, such as models in which the addition operation is not commutative, or models in which, for example,  $10=0$ . We restrict the class of models of a MEL specification only to those which are *initial*, that is, those which obey the *no junk no confusion* principle; therefore, our specifications have *initial semantics* [49] in this paper. Intuitively, that means that only those models are allowed in which all elements can be “constructed” from smaller elements and in which no terms which cannot be proved equal are interpreted to the same elements.

By reducing the class of models, one can enlarge the class of sound inference rules. A major benefit one gets under initial semantics is that *proofs by induction become valid*. Since the proof of correctness for the main algorithm in this paper is done by induction on the structure of the temporal formula to monitor, it is important for the reader to be aware that the specifications presented from now on have initial semantics.

### Syntactic Sugar Conventions

To make specifications easier to read, the following syntactic sugar conventions are widely accepted:

**Subsorts.** Given sorts  $s, s'$  with  $\pi(s) = \pi(s') = k$ , the declaration  $s < s'$  is syntactic sugar for the conditional membership  $(\forall x : k) x : s' \text{ if } x : s$ .

**Operations.** If  $\sigma \in \Omega_{k_1 \dots k_n, k}$  and  $s_1, \dots, s_n, s \in S$  with  $\pi(s_1) = k_1, \dots, \pi(s_n) = k_n, \pi(s) = k$ , then the declaration  $\sigma : s_1 \cdots s_n \rightarrow s$  is syntactic sugar for  $(\forall x_1 : k_1, \dots, x_n : k_n) \sigma(x_1, \dots, x_n) : s \text{ if } x_1 : s_1 \wedge \dots \wedge x_n : s_n$ .

**Variables.**  $(\forall x : s, X) \text{ } a \text{ if } a_1 \wedge \dots \wedge a_n$  is syntactic sugar for the  $\Omega$ -sentence  $(\forall x : \pi(s), X) \text{ } a \text{ if } a_1 \wedge \dots \wedge a_n \wedge x : s$ . With this, the operation declaration  $\sigma : s_1 \cdots s_n \rightarrow s$  above is equivalent to  $(\forall x_1 : s_1, \dots, x_n : s_n) \sigma(x_1, \dots, x_n) : s$ .

### 2.1.2 Maude

Maude [24] is a freely distributed high-performance system in the OBJ [50] algebraic specification family, supporting both rewriting logic [93] and membership equational logic [94]. Because of its efficient rewriting engine, able to execute 3 million rewriting steps per second on standard PCs, and because of its metalanguage features, Maude turns out to be an excellent tool to create executable environments for various logics, models of computation, theorem provers, and even programming languages. We were delighted to notice how easily we could implement and efficiently validate our algorithms for testing LTL formulae on finite event traces in Maude, admittedly a tedious task in C++ or Java, and hence decided to use Maude at least for the prototyping stage of our runtime check algorithms.

We informally describe some of Maude’s features via examples in this section, referring the interested reader to its manual [24] for more details. The examples discussed in this subsection are not random. On the one hand they show all the major features of Maude that we need, while on the other hand they are part of our current JPaX implementation; several references to them will be made later in the paper. Maude supports modularization in the OBJ style. There are various kinds of modules, but we use only functional modules which follow the pattern “`fmod <name> is <body> endfm`”, and which have initial semantics. The body of a functional module consists of a collection of declarations, of which we are using importation, sorts, subsorts, operations, variables and equations, usually in this order.

### Defining Logics for Monitoring

In the following we introduce some modules that we think are general enough to be used within any logical environment for program monitoring that one would want to implement by rewriting. The first one simply defines atomic propositions as an abstract data type having one sort, `Atom`, and no operations or constraints:

```
fmod ATOM is
  sort Atom .
```

```
endfm
```

The actual names of atomic propositions will be automatically generated in another module that extends `ATOM`, as constants of sort `Atom`. These will be generated by the observer at the initialization of monitoring, from the actual properties that one wants to monitor.

An important concept in program monitoring is that of an (abstract) execution trace, which consists of a finite list of events. We abstract a single event by a list of atoms, those that hold after the action that generated the event took place. The values of the atomic propositions are updated by the observer according to the actual state of the executing program and then sent to Maude as a term of sort `Event` (more details regarding the communication between the running program and Maude will be given later):

```
fmod TRACE is
  protecting ATOM .
  sorts Event Event* Trace .
  subsorts Atom < Event < Event* < Trace .
  op empty : -> Event .
  op _ : Event Event -> Event [assoc comm id: empty prec 23] .
  var A : Atom .
  eq A A = A .
  op _* : Event -> Event* .
  op _ , _ : Event Trace -> Trace [prec 25] .
endfm
```

The statement `protecting ATOM` imports the module `ATOM` without changing its initial semantics. The above is a compact way to use *mix-fix*<sup>1</sup> and order-sorted notation to define an abstract data type of traces: a trace is a comma separated list of events, where an event is itself a *set* of atoms. The `subsorts` declaration declares `Atom` to be a subsort of `Event`, which in turn is a subsort of `Event*` which is a subsort of `Trace`. Since elements of a subsort can occur as elements of a supersort without explicit lifting, we have as a consequence that a single event is also a trace, consisting of one event. Likewise, an atomic proposition can occur as an event, containing only this atomic proposition.

Operations can have attributes, such as associativity (A), commutativity (C), identity (I) as well as precedences, which are written between square brackets. When a binary operation is declared using the attributes A, C, and/or I, Maude uses built-in efficient specialized algorithms for matching

---

<sup>1</sup>Underscores are places for arguments.

and rewriting. However, semantically speaking, the A, C, and/or I attributes can be replaced by their corresponding equations. The attribute `prec` gives a precedence to an operator<sup>2</sup>, thus eliminating the need for most parentheses. Notice the special sort `Event*` which stands for terminal events, i.e., events that occur at the end of traces. Any event can potentially occur at the end of a trace. It is often the case that ending events are treated differently, like in the case of finite trace linear temporal logic; for this reason, we have introduced the operation `_*` which marks an event as terminal.

An event is defined as a set of atoms which should in fact be thought of as the set of all those atoms which “hold” in the new state of the event emitting program. Note the idempotency equation “`eq A A = A`”, which ensures that an event is indeed a set. On the other hand, a trace is an ordered list of events which can potentially have repetitions of events. For example, the event “ $x = 5$ ” can occur several times during the execution of a program. Note that there is no need and consequently no definition of an empty trace.

Syntax and semantics are basic requirements to any logic. The following module introduces what we believe are the basic ingredients of monitoring logics, i.e., logics used for specifying monitoring requirements:

```
fmod LOGICS-BASIC is
  protecting TRACE .
  sort Formula .
  subsort Atom < Formula .
  ops true false : -> Formula .
  op [_] : Formula -> Bool .
  eq [true] = true .
  eq [false] = false .

  var A : Atom .
  var T : Trace .
  var E : Event .
  var E* : Event* .
  op _{[_]} : Formula Event* -> Formula [prec 10] .
  eq true{E*} = true .
  eq false{E*} = false .
  eq A{A E} = true .
  eq A{E} = false [owise] .
  eq A{E *} = A{E} .
```

---

<sup>2</sup>The lower the precedence number, the tighter the binding.

```

op _|=_ : Trace Formula -> Bool [prec 30] .
eq T |= true  = true .
eq T |= false = false .
eq E   |= A   = [A{E}] .
eq E,T |= A   = E |= A .
endfm

```

The first block of declarations introduces the sort `Formula` which can be thought of as a generic sort for any well-formed formula in any logic. There are two designated formulae, namely `true` and `false`, with the obvious meaning in any monitoring logic. The sort `Bool` is built-in in Maude together with two constants `true` and `false`, which are different from those of sort `Formula`, and a generic operator `if_then_else-fi`. The “interpretation” operator `[_]` maps a formula to a Boolean value. Each logic implemented on top of LOGICS-BASIC is free to define it appropriately; here we only give the obvious mappings of `true` and `false` of `Formula` into `true` and `false` of `Bool`.

The second block defines the operation `_{}_` which takes a formula and an event and yields another formula. The intuition for this operation is that it “evaluates” the formula in the new state and produces a proof obligation as another formula for the subsequent events. If the returned formula is `true` or `false` then it means that the formula was satisfied or violated, regardless of the rest of the execution trace; in this case, a message can be returned by the observer. Each logic will further complete the definition of this operator. Note that the equation “`eq A{A E} = true`” speculates Maude’s capability of performing matching modulo associativity, commutativity and identity (the attributes of the *set* concatenation on events); it basically says that `A{E}` is `true` if `E` contains the atom `A`. The next equation contains the attribute `[owise]`, stating that it should be applied only if any other equation fails to apply at a particular position.

Finally, the satisfaction relation is defined. Two obvious equations deal with the formulae `true` and `false`. The last two equations state that a trace satisfies an atomic proposition `A` if evaluating that atomic proposition `A` on the first event in the trace yields `true`. The remaining elements in the trace do not matter because `A` is a simple atom, so it refers to only the current state.

## Defining Propositional Calculus

A rewriting decision procedure for propositional calculus due to Hsiang [71] is adapted and presented. It provides the usual connectives `_/_` and `_/\_`.

(and),  $\_++\_$  (exclusive or),  $\_\\/_$  (or),  $\!_$  (negation),  $\_\\_$  (implication), and  $\_<->_$  (equivalence). The procedure reduces tautological formulae to the constant `true` and all the others to some canonical form modulo associativity and commutativity. An unusual aspect of this procedure is that a canonical form consists of an exclusive or of conjunctions. In fact, this choice of basic operators corresponds to regarding propositional calculus as a Boolean ring rather than as a Boolean algebra. A major advantage of this choice is that normal forms are unique modulo associativity and commutativity. Even if propositional calculus is very basic to almost any logical environment, we decided to keep it as a separate logic instead of being part of the logic infrastructure of JPAX. One reason for this decision is that its operational semantics could be in conflict with other logics, for example ones in which conjunctive normal forms are desired.

An OBJ3 code for this procedure appeared in [50]. Below we give its obvious translation to Maude together with its finite trace semantics, noticing that Hsiang [71] showed that this rewriting system modulo associativity and commutativity is Church-Rosser and terminates. The Maude team was probably also inspired by this procedure, since the builtin `BOOL` module is very similar.

```
fmod PROP-CALC is
  extending LOGICS-BASIC .
*** Constructors ***
  op _/__ : Formula Formula -> Formula [assoc comm prec 15] .
  op _++_ : Formula Formula -> Formula [assoc comm prec 17] .
  vars X Y Z : Formula .
  eq true /\ X = X .
  eq false /\ X = false .
  eq X /\ X = X .
  eq false ++ X = X .
  eq X ++ X = false .
  eq X /\ (Y ++ Z) = X /\ Y ++ X /\ Z .
*** Derived operators ***
  op _\\/_ : Formula Formula -> Formula [assoc prec 19] .
  op !_ : Formula -> Formula [prec 13] .
  op _->_ : Formula Formula -> Formula [prec 21] .
  op _<->_ : Formula Formula -> Formula [prec 23] .
  eq X \\ Y = X /\ Y ++ X ++ Y .
  eq ! X = true ++ X .
  eq X -> Y = true ++ X ++ X /\ Y .
  eq X <-> Y = true ++ X ++ Y .
*** Finite trace semantics
```

```

var T : Trace .
var E* : Event* .
eq T |= X /\ Y = T |= X and T |= Y .
eq T |= X ++ Y = T |= X xor T |= Y .
eq (X /\ Y){E*} = X{E*} /\ Y{E*} .
eq (X ++ Y){E*} = X{E*} ++ Y{E*} .
eq [X /\ Y] = [X] and [Y] .
eq [X ++ Y] = [X] xor [Y] .
endfm

```

The statement “`extending LOGICS-BASIC`” imports the module `LOGICS-BASIC` with the reserve that its initial semantics can be extended. The operators “`and`” and “`xor`” come from the Maude’s built-in module `BOOL` which is automatically imported by any other module.

Operators are declared with special attributes, such as `assoc` and `comm`, which enable Maude to use its specialized efficient internal rewriting algorithms. Once the module above is loaded<sup>3</sup> in Maude, reductions can be done as follows:

```

reduce a -> b /\ c <-> (a -> b) /\ (a -> c) . ***> should be true
reduce a <-> ! b . ***> should be a ++ b

```

Notice that one should first declare the constants `a`, `b` and `c`. The last six equations in the module `PROP-CALC` are related to the semantics of propositional calculus. The default evaluation strategy for `[_]` is eager, so `[X]` will first evaluate `X` using propositional calculus reasoning and then will apply one of the last two equations if needed; these equations will not be applied normally in practical reductions, they are useful only in the correctness proof stated by Theorem 16.

---

*end from J.ASE’05*

---



---

<sup>3</sup>Either by typing it or using the command “`in <filename>`”.



## Chapter 3

# Safety Properties

---

*safety property = every violation occurs after a finite execution.*

---

---

*later, when talking about LTL, discuss the classification of safety properties in [84]*

---

---

*In the literature, what we call “prefixes” are also called “good prefixes”, while the rest of the prefixes are called “bad prefixes”.*

---

Intuitively, a safety property of a system is one stating that the system cannot “go wrong”, or, as Lamport [85] put it, that the “bad thing” never happens. In other words, in order for a system to violate a safety property, it should eventually “go wrong” or the “bad thing” should eventually happen. There is a very strong relationship between safety properties and runtime monitoring: if a safety property is violated by a running system, then the violation should happen *during* the execution of the system, in a finite amount of time, so a monitor for that property observing the running system should be able to detect the violation; an additional point in the favor of monitoring is that, if a system violates a safety property at some moment during its execution, then there is no way for the system to continue its execution to eventually satisfy the property, so a monitor needs not wait for a better future once it detects a bad present/past.

State properties or assertions that need only the current state of the running system to check whether they are violated or not, such as “no division by 0”, or “ $x$  positive”, or no deadlock, are common safety properties;

once violated, one can stop the computation or take corrective measures. However, there are also interesting safety properties that involve more than one state of the system, such as “if one uses resource  $x$  then one must have authenticated at some moment in the past”, or “any start of a process must be followed by a stop within 10 units of time”, or “take command from user only if the user has logged in at some moment in the past and has not logged out since then”, etc. Needless to say that the atomic events, or states, which form execution traces on which safety properties are defined, can be quite abstract: not all the details of a system execution are relevant for the particular safety property of interest. In the context of monitoring, these relevant events or states can be extracted by means of appropriate instrumentation of the system. For example, runtime monitoring systems such as Tracematches [4] and MOP [20] use aspect-oriented technology to “hook” relevant observation points and appropriate event filters in a system.

It is customary to define safety properties as properties over *infinite traces*, to capture the intuition that they are defined for systems that can potentially run forever, such as reactive systems. A point in favor of infinite traces is that finite traces can be regarded as special cases of infinite traces, namely ones that “stutter” indefinitely in their last state (see, for example, Abadi and Lamport [2, 3]). Infinite traces are particularly desirable when one specifies safety properties using formalisms that have infinite-trace semantics, such as linear temporal logics or corresponding automata.

While “infinity” is a convenient abstraction that is relatively broadly-accepted nowadays in mathematics and in theoretical foundations of computer science, there is no evidence so far that a system can have an infinite-trace behavior (we have not seen any). A disclaimer is in place here: we do *not* advocate finite-traces as a foundation for safety properties; all we try to do is to argue that, just because they can be seen as a special case of infinite traces, finite traces are not entirely uninteresting. For example, a safety property associated to a one-time-access key issued to a client can be “activate, then use at most once, then close”. Using regular patterns over the alphabet of relevant events  $\Sigma = \{\text{activate}, \text{use}, \text{close}\}$ , this safety property can be expressed as “ $\text{activate} \cdot (\epsilon + \text{use}) \cdot \text{close}$ ”; any trace that is not a prefix of the language of this regular expression violates the property, including any other activation or use of the key after it was closed. While these finite-trace safety properties can easily be expressed as infinite-trace safety properties, we believe that that would be more artificial than simply accepting that in practice we deal with many finite-trace safety properties.

In this section we discuss various approaches to formalize safety properties and show that they are ultimately directly or indirectly equivalent. We categorize them into finite-trace safety properties, infinite-trace safety properties, and finite- and infinite-trace safety properties:

1. Section 3.1 defines safety properties over finite traces as prefix closed properties. A subset of finite-trace safety properties, that we call *persistent*, contain only traces that “have a future” within the property, that is, finite traces that can be continued into other finite traces that are also in the safety property. Persistent safety properties appear to be the right finite-trace variant that corresponds faithfully to the more conventional infinite-trace safety properties. Even though persistent safety properties form a proper subset of finite-trace safety properties and each finite-trace safety property has a largest persistent safety property included in it, we show that there is in fact a bijection between safety properties and persistent safety properties by showing them both to have the cardinal of the continuum  $c$ .
2. In Section 3.2, we consider two standard infinite-trace definitions of a safety property, one based on the intuition that violating behaviors must manifest so after a finite number of events and the other based on the intuition of a safety property as a closed set in an appropriate topology over infinite-traces. We show them both equivalent to persistent safety properties over finite traces, by constructing an explicit bijection (as opposed to using cardinality arguments and infer the existence of a bijection); consequently, infinite-trace safety properties also have the cardinal of the continuum  $c$ . Since closed sets of real numbers are in a bijective correspondence with the real numbers, we indirectly rediscover Alpern and Schneider’s result [5] stating that infinite-trace safety properties correspond to closed sets in infinite-trace topology.
3. Section 3.3 considers safety properties defined over both finite and infinite traces. We discuss two definitions of such safety properties encountered in the literature, and, using cardinality arguments, we show their equivalence with safety properties over only finite traces. In particular, safety properties over finite and infinite traces also have the cardinality of the continuum  $c$ . We also show that prefix-closeness is not a sufficient condition to characterize (not even bijectively) such safety properties, by showing that there are significantly more ( $2^c$ ) prefix-closed properties over finite and infinite traces than safety properties.

Therefore, each of the classes of safety properties is in bijection with the real numbers. Since there are so many safety properties, we can also insightfully conclude that there is *no* enumerable mechanism to define all the safety properties, because  $\aleph_0 \leq c$ . Therefore, particular logical or syntactic recursive formalisms can only define *some* of the safety properties, but not all of them.

### 3.1 Finite Traces

One of the most common intuitions for a safety property is as a prefix-closed set of finite traces. This captures best the intuition that once something bad happened, there is no way to recover: if  $w \notin P$  then there is no  $u$  such that  $P(wu)$ , which is equivalent to saying that if  $P(wu)$  then  $P(w)$ , which is equivalent to saying that  $P$  is prefix closed. From a monitoring perspective, a prefix closed property can be regarded as one containing all the good (complete or partial) behaviors of the observed system: once a state is encountered that does not form a good behavior together with the previously observed states, then a violation can be reported.

**Definition 1** Let  $\text{prefixes} : \Sigma^* \rightarrow \mathcal{P}(\Sigma^*)$  be the prefix function returning for any finite trace all its prefixes, and let  $\text{prefixes} : \mathcal{P}(\Sigma^*) \rightarrow \mathcal{P}(\Sigma^*)$  be its corresponding closure operator that takes sets of finite traces and closes them under prefixes.

Note that  $\text{prefixes} : \mathcal{P}(\Sigma^*) \rightarrow \mathcal{P}(\Sigma^*)$  is indeed a closure operator (Exercise 2): it is extensive ( $P \subseteq \text{prefixes}(P)$ ), monotone ( $P \subseteq P'$  implies  $\text{prefixes}(P) \subseteq \text{prefixes}(P')$ ), and idempotent ( $\text{prefixes}(\text{prefixes}(P)) = \text{prefixes}(P)$ ).

**Definition 2** Let  $\text{Safety}^*$  be the set of finite-trace prefix-closed properties, that is, the set  $\{P \in \mathcal{P}(\Sigma^*) \mid P = \text{prefixes}(P)\}$ . In other words,  $\text{Safety}^*$  is the set of fixed points of the prefix operator  $\text{prefixes} : \mathcal{P}(\Sigma^*) \rightarrow \mathcal{P}(\Sigma^*)$ .

The star superscript in  $\text{Safety}^*$  reflects that its traces are finite; in the next section we will define a set  $\text{Safety}^\omega$  of infinite-trace safety properties. Since  $\text{prefixes}(P) \in \text{Safety}^*$  for any  $P \in \mathcal{P}(\Sigma^*)$ , we can assume from here on that  $\text{prefixes} : \mathcal{P}(\Sigma^*) \rightarrow \mathcal{P}(\Sigma^*)$  is actually a function  $\mathcal{P}(\Sigma^*) \rightarrow \text{Safety}^*$ .

**Example 1** Consider the one-time-access key safety property discussed above, saying that a client can “activate, then use at most once, and then

close” the key. If  $\Sigma = \{\text{activate}, \text{use}, \text{close}\}$ , then this safety property can be expressed as the finite set of finite words

$$\{\epsilon, \text{activate}, \text{activate close}, \text{activate use}, \text{activate use close}\}$$

No other behavior is allowed. Now suppose that the safety policy is extended to allow multiple uses of the key once activated, but still no further events once it is closed. The extended safety property has now infinitely many finite-traces:

$$\{\epsilon\} \cup \{\text{activate}\} \cdot \{\text{use}^n \mid n \in \mathbb{N}\} \cdot \{\epsilon, \text{close}\}.$$

Note that this property is indeed prefix-closed. A monitor in charge of online checking this safety property would report a violation if the first event is not *activate*, or if it encounters any second *activate* event, or if it encounters any event after a *close* event is observed, including another *close* event.

It is interesting to note that this finite-trace safety property encompasses both finite and infinite aspects. For example, it does not preclude behaviors in which one sees an *activate* event and then an arbitrary number of *use* events; *use* events can persist indefinitely after an *activate* event without violating the property. On the other hand, once a *close* event is encountered, no other event can be further seen. We will shortly see that the safety property above properly includes the *persistent* safety property  $\{\epsilon\} \cup \{\text{activate use}^n \mid n \in \mathbb{N}\}$ , which corresponds to the infinite-trace safety property  $\{\text{activate use}^\omega\}$ .  $\square$

While prefix closeness seems to be the right requirement for a safety property, one can argue that it is not sufficient. For example, in the context of reactive systems that supposedly run forever, one may think of a safety property as one containing safe finite traces, that is, ones for which the reactive system can always find a way to continue its execution safely. The definition of safety properties above includes, among other safety properties, the empty set of traces as well as all prefix-closed *finite* sets of finite traces; any reactive system will eventually violate such safety properties, so one can say that the definition of safety property above is too generous.

We next define *persistent safety properties* as ones that always allow a future; intuitively, an observed reactive system that is in a safe state can always (if persistent enough) find a way to continue its execution to a next safe state. This notion is reminiscent of “feasibility”, a semantic characterization of fairness in [8], and of “machine closeness” [2, 112], also used in the context of fairness.

**Definition 3** Let  $\text{PersistentSafety}^*$  be the set of finite-trace persistent safety properties, that is, safety properties  $P \in \text{Safety}^*$  such that if  $P(w)$  for some  $w \in \Sigma^*$  then there is some  $a \in \Sigma$  such that  $P(wa)$ .

If a persistent safety property is non-empty, then note that it must contain an infinite number of words. The persistency aspect of a finite-trace safety property can be regarded, in some sense, as a liveness argument. Indeed, assuming that it is a “good thing” for a trace to be indefinitely continued, then a persistent safety property is one in which the “good thing” always eventually happens. If one takes the liberty to regard “stuck” computations as unfair, then the persistency aspect above can also be regarded as a fairness argument.

Another way to think of persistent safety properties is as a means to refer to infinite behaviors by means of finite traces. This view is, in some sense, dual to the more common approach to regard finite behaviors as infinite behaviors that stutter infinitely in a “last” state (see, for example, Abadi and Lamport [2, 3] for a formalization of such last-state infinite stuttering).

Note that if  $\Sigma$  is a degenerate set of events containing only one element, that is, if  $|\Sigma| = 1$ , then  $|\text{Safety}^*| = \aleph_0$  and  $|\text{PersistentSafety}^*| = 2$ ; indeed, if  $\Sigma = \{a\}$  then  $\text{Safety}^*$  contains precisely the finite properties  $a^{\leq n} = \{a^i \mid 0 \leq i \leq n\}$  for each  $n \in \mathbb{N}$  plus the infinite property  $\{a^n \mid n \in \mathbb{N}\}$ , so a total of  $\aleph_0 + 1 = \aleph_0$  properties, while  $\text{PersistentSafety}^*$  contains only two properties, namely  $\emptyset$  and  $\{a^n \mid n \in \mathbb{N}\}$ . The case when there is only one event or state in  $\Sigma$  is neither interesting nor practical. Therefore, from here on in this paper we take the liberty to assume that  $|\Sigma| \geq 2$ . Since in practice  $\Sigma$  contains states or events generated by a computer, for simplicity in stating some of the subsequent results, we also take the liberty to assume that  $|\Sigma| \leq \aleph_0$ ; therefore,  $\Sigma$  can be any finite or recursively enumerable set, including  $\mathbb{N}$ ,  $\mathbb{N}_\infty$ ,  $\mathbb{Q}$ , etc., but cannot be  $\mathbb{R}$  or any set “larger” than  $\mathbb{R}$ . With these assumptions, it follows that  $|\Sigma^*| = \aleph_0$  (finite words are recursively enumerable) and  $|\Sigma^\omega| = c$  (infinite streams have the cardinality of the continuum).

**Proposition 1**  $\text{Safety}^*$  and  $\text{PersistentSafety}^*$  are closed under union;  $\text{Safety}^*$  is also closed under intersection.

**Proof:** The union and the intersection of prefix-closed properties is also prefix-closed. Also, the union of persistent prefix-closed properties is also persistent.  $\square$

The intersection of persistent safety properties may not be persistent:

**Example 2** Let  $\Sigma$  be the set  $\{0, 1\}$ . Let  $P = \{1^m \mid m \in \mathbb{N}\}$  and  $P' = \{\epsilon\} \cup \{10^m \mid m \in \mathbb{N}\}$  be two persistent safety properties, where  $\epsilon$  is the empty word (the word containing no letters). Then  $P \cap P'$  is the finite safety property  $\{\epsilon, 1\}$ , which is not persistent. If one thinks that this happened because  $P \cap P'$  does not contain any proper (i.e., non-empty) persistent property, then one can take instead the persistent safety properties  $P = \{0^n \mid n \in \mathbb{N}\} \cdot \{1^m \mid m \in \mathbb{N}\}$  and  $P' = \{0^n \mid n \in \mathbb{N}\} \cdot (\{\epsilon\} \cup \{10^m \mid m \in \mathbb{N}\})$ , whose intersection is the safety property  $\{0^n \mid n \in \mathbb{N}\} \cup \{0^n 1 \mid n \in \mathbb{N}\}$ . This safety property is not persistent because its words ending in 1 cannot persist, but it contains the proper persistent safety property  $\{0^n \mid n \in \mathbb{N}\}$ .  $\square$

Therefore, we can associate to any safety property in  $\text{Safety}^*$  a largest persistent safety property in  $\text{PersistentSafety}^*$ , by simply taking the union of all persistent safety properties that are included in the original safety property (the empty property is one of them, the smallest):

**Definition 4** For a safety property  $P \in \text{Safety}^*$ , let  $P^\circ \in \text{PersistentSafety}^*$  be the largest persistent safety property with  $P^\circ \subseteq P$ .

The following example shows that one may need to eliminate infinitely many words from a safety property in order to obtain a persistent safety property:

**Example 3** Let  $\Sigma = \{0, 1\}$  and let  $P$  be the safety property  $\{0^n \mid n \in \mathbb{N}\} \cup \{0^n 1 \mid n \in \mathbb{N}\}$ . Then  $P^\circ$  can contain no word ending with a 1 and can contain all the words of 0's. Therefore,  $P^\circ = \{0^n \mid n \in \mathbb{N}\}$ .  $\square$

Finite safety properties obviously cannot contain any non-empty persistent safety property, that is,  $P^\circ = \emptyset$  if  $P$  is finite. But what if  $P$  is infinite? Is it always the case that it contains a non-empty persistent safety property? Interestingly, it turns out that this is true if and only if  $\Sigma$  is finite:

**Proposition 2** If  $\Sigma$  is finite and  $P$  is a safety property containing infinitely many words, then  $P^\circ \neq \emptyset$ .

**Proof:** For each letter  $a \in \Sigma$ , let us define the *derivative of  $P$  wrt  $a$* , written  $\delta_a(P)$ , as the language  $\{w \in \Sigma^* \mid aw \in P\}$ . Since

$$P = \{\epsilon\} \cup \bigcup_{a \in \Sigma} \{a\} \cdot \delta_a(P)$$

since  $\Sigma$  is finite, and since  $P$  is infinite, it follows that there is some  $a_1 \in \Sigma$  such that  $\delta_{a_1}(P)$  is infinite; note that  $a_1 \in P$  since  $P$  is prefix closed. Similarly, since  $\delta_{a_1}(P)$  is infinite, there is some  $a_2 \in \Sigma$  such that  $\delta_{a_2}(\delta_{a_1}(P))$  is infinite and  $a_1a_2 \in P$ . Iterating this reasoning, we can find some  $a_n \in \Sigma$  for each  $n \in \mathbb{N}$ , such that  $a_1a_2 \dots a_n \in P$  and  $\delta_{a_n}(\dots(\delta_{a_2}(\delta_{a_1}(P)))\dots)$  is infinite, that is, the set  $\{w \in \Sigma^* \mid a_1a_2 \dots a_n w \in P\}$  is infinite. It is now easy to see that the set  $\{a_1a_2 \dots a_n \mid n \in \mathbb{N}\} \subseteq P$  is persistent. Therefore,  $P^\circ \neq \emptyset$ .  $\square$

The following example shows that  $\Sigma$  must indeed be finite in order for the result above to hold:

**Example 4** Consider some infinite set of events or states  $\Sigma$ . Then we can label distinct elements in  $\Sigma$  with distinct labels in  $\mathbb{N} \cup \{\infty\}$ . We only need these elements from  $\Sigma$ ; therefore, without loss of generality, we can assume that  $\Sigma = \mathbb{N} \cup \{\infty\}$ . Let  $P$  be the safety property

$$\{\epsilon\} \cup \{\infty n(n-1) \dots (m+1)m \mid 0 \leq m \leq n+1\},$$

where  $\epsilon$  is the empty word (the word containing no letters) and  $n \dots (n+1)$  is also the empty word for any  $n \in \mathbb{N}$ . Then  $P^\circ$  is the empty property. Indeed, note that any persistent safety property  $P'$  included in  $P$  cannot have traces ending in 0, because those cannot be continued into other traces in  $P$ ; since  $P'$  cannot contain traces ending in 0, it cannot contain traces ending in 1 either, because such traces can only be continued with a 0 letter into traces in  $P$ , but those traces have already been decided that cannot be part of  $P'$ ; inductively, one can show that  $P'$  can contain no words ending in letters that are natural numbers in  $\mathbb{N}$ . Since the only trace in  $P$  ending in  $\infty$  is  $\infty$  itself and since  $\infty$  can only be continued with a natural number letter into a trace in  $P$  but such trace cannot belong to  $P'$ , we deduce that  $P'$  can contain no word with letters in  $\Sigma$ . In particular,  $P^\circ$  must be empty.  $\square$

Even though we know that the largest persistent safety property  $P^\circ$  included into a safety property  $P$  always exists because **PersistentSafety**<sup>\*</sup> is closed under union, we would like to have a more constructive way to obtain it. A first and obvious thing to do is to eliminate from  $P$  all the “stuck” computations, that is, those which cannot be added any new state to obtain a trace that is also in  $P$ . This removal step does not destroy the prefix-closeness of  $P$ , but it may reveal new computations which are stuck. By iteratively eliminating all the computations that get stuck in a finite number of steps, one would expect to obtain a persistent safety property,

namely precisely  $P^\circ$ . It turns out that this is indeed true only if  $\Sigma$  is finite. If that is the case, then the following can also be used as an alternative definition of  $P^\circ$ :

**Proposition 3** *Given safety property  $P \in \mathbf{Safety}^*$ , then let  $P^-$  be the property  $\{w \in P \mid (\exists a \in \Sigma) wa \in P\}$ . Also, let  $\{P_i \mid i \in \mathbb{N}\}$  be properties defined as  $P_0 = P$  and  $P_{i+1} = P_i^-$  for all  $i \geq 0$ . Then  $P^\circ = \bigcap_{i \geq 0} P_i$  whenever  $\Sigma$  is finite.*

**Proof:** It is easy to see that if  $P$  is prefix-closed then  $P^- \subseteq P$  is also prefix-closed, so  $P^-$  is also a property in  $\mathbf{Safety}^*$ . Therefore, the properties  $P_i$  form a sequence  $P = P_0 \supseteq P_1 \supseteq P_2 \supseteq \dots$  of increasingly smaller safety properties.

Let us first prove that  $\bigcap_{i \geq 0} P_i$  is a persistent safety property. Assume by contradiction that for some  $w \in \bigcap_{i \geq 0} P_i$  there is no  $a \in \Sigma$  such that  $wa \in \bigcap_{i \geq 0} P_i$ . In other words, we can find for each  $a \in \Sigma$  some  $i_a \geq 0$  such that  $wa \notin P_{i_a}$ . Since  $\Sigma$  is finite, we can let  $i$  be the largest among the natural numbers  $i_a \in \mathbb{N}$  for all  $a \in \Sigma$ . Since  $P_i \subseteq P_{i_a}$  for all  $a \in \Sigma$ , it should be clear that there is no  $a \in \Sigma$  such that  $wa \in P_i$ , which means that  $w \notin P_{i+1}$ . This contradicts the fact that  $w \in \bigcap_{i \geq 0} P_i$ . Therefore,  $\bigcap_{i \geq 0} P_i \in \mathbf{PersistentSafety}^*$ .

Let us now prove that  $\bigcap_{i \geq 0} P_i$  is the largest persistent safety property included in  $P$ . Let  $P'$  be any persistent safety property included in  $P$ . We show by induction on  $i$  that  $P' \subseteq P_i$  for all  $i \in \mathbb{N}$ . The base case,  $P' \subseteq P_0$ , is obvious. Suppose that  $P' \subseteq P_i$  for some  $i \in \mathbb{N}$  and let  $w \in P'$ . Since  $P'$  is persistent, there is some  $a \in \Sigma$  such that  $wa \in P' \subseteq P_i$ , which means that  $w \in P_{i+1}$ . Since  $w$  was chosen arbitrarily, it follows that  $P' \subseteq P_{i+1}$ . Therefore,  $P' \subseteq \bigcap_{i \geq 0} P_i$ .  $\square$

We next show that the finiteness of  $\Sigma$  was a necessary requirement in order for the result above to hold. In other words, we show that if  $\Sigma$  is allowed to be infinite then we can find a safety property  $P \in \mathbf{Safety}^*$  over  $\Sigma$  such that  $P^\circ \in \mathbf{PersistentSafety}^*$  and  $\bigcap_{i \geq 0} P_i \in \mathbf{Safety}^*$  are distinct. Since we showed in the proof of Proposition 3 that any persistent safety property  $P'$  is included in  $\bigcap_{i \geq 0} P_i$ , it follows that  $P^\circ \subseteq \bigcap_{i \geq 0} P_i$ . Since  $P^\circ$  is the largest persistent safety property included in  $P$ , one can easily show that  $P^\circ = (\bigcap_{i \geq 0} P_i)^\circ$ . Therefore, it suffices to find a safety property  $P$  such that  $\bigcap_{i \geq 0} P_i$  is not persistent, which is what we do in the next example:

**Example 5** Consider the safety property  $P$  over infinite  $\Sigma = \mathbb{N} \cup \{\infty\}$  discussed in Example 4, namely  $\{\epsilon\} \cup \{\infty n(n-1) \dots (m+1)m \mid 0 \leq m \leq$

$n + 1\}$ . Then one can easily show by induction on  $i \in \mathbb{N}$  that the properties  $P_i$  defined in Proposition 3 are the sets  $\{\epsilon\} \cup \{\infty n(n-1) \dots (m+1)m \mid i \leq m \leq n+1\}$ ; in other words, each  $P_i$  excludes from  $P$  all the words whose last letters are smaller than  $i$  when regarded as natural numbers. Then the intersection  $\bigcap_{i \geq 0} P_i$  contains no trace ending in a natural number; the only possibility left is then  $\bigcap_{i \geq 0} P_i = \{\epsilon, \infty\}$ , which is different from  $P^\circ = \emptyset$  (see Example 4).

One may argue that  $P^\circ \neq \bigcap_{i \geq 0} P_i$  above happened precisely because  $P^\circ$  was empty. One can instead pick the safety property  $Q = \{0^n \mid n \in \mathbb{N}\} \cdot P$ . Then one can show following the same idea as in Example 4 that  $Q^\circ = \{0^n \mid n \in \mathbb{N}\}$ . Further, one can show that  $Q_i = \{0^n \mid n \in \mathbb{N}\} \cdot P_i$ , so  $\bigcap_{i \geq 0} Q_i = \{0^n \mid n \in \mathbb{N}\} \cup \{0^n \infty \mid n \in \mathbb{N}\}$ , which is different from  $Q^\circ$ .  $\square$

Persistency is reminiscent of “feasibility” introduced by Apt et al. [8] in the context of fairness, and of “machine closeness” introduced by Abadi and Lamport [2, 3] (see also Schneider [112]) in the context of refinement. Let us use the terminology “machine closeness”: a property  $L$  (typically a liveness or a fairness property) is *machine closed* for a property  $M$  (typically given as the language of some state machine) iff  $L$  does not prohibit any of the observable runtime behaviors of  $M$ , that is, iff  $\text{prefixes}(M) = \text{prefixes}(M \cap L)$ ; for example, if  $M$  is the total property (i.e., every event is possible at any moment, i.e.,  $M = \Sigma^*$ ) and  $L$  is the property stating that “always eventually event  $a$ ”, then any prefix of  $M$  can be continued to obtain a property satisfying  $L$ . Persistency is related to machine closeness in that a safety property  $P$  is persistent if and only if  $P^\circ$  is machine closed for  $P$ . In other words, there is nothing  $P$  can do in a finite amount of time that  $P^\circ$  cannot do. However, there is a caveat here: since liveness and fairness are inherently infinite-trace notions, machine closeness (or feasibility) have been introduced in the context of infinite-traces. On the other hand, persistency makes sense only in the context of finite traces.

It is clear that  $\text{PersistentSafety}^*$  is properly included in  $\text{Safety}^*$ . Yet, we next show that, surprisingly, there is a bijective correspondence between  $\text{Safety}^*$  and  $\text{PersistentSafety}^*$ , both having the cardinal of the continuum:

**Theorem 2**  $|\text{PersistentSafety}^*| = |\text{Safety}^*| = c$ .

**Proof:** Since  $\Sigma^*$  is recursively enumerable and since  $2^{\aleph_0} = c$ , we can readily infer that  $|\text{PersistentSafety}^*| \leq |\text{Safety}^*| \leq |\mathcal{P}(\Sigma^*)| = c$ .

Let us now define an injective function  $\varphi$  from the open interval of real numbers  $(0, 1)$  to  $\text{PersistentSafety}^*$ . Since  $|\Sigma| \geq 2$ , let us distinguish

two different elements in  $\Sigma$  and let us label them  $\bar{0}$  and  $\bar{1}$ . For a real  $r \in (0, 1)$ , let  $\varphi(r)$  be the set  $\{\bar{\alpha} \mid \alpha \in \{0, 1\}^* \text{ and } 0.\alpha < r\}$ , where  $0.\alpha$  is the (rational) number in  $(0, 1)$  whose decimals in binary representation are  $\alpha$ , and where  $\bar{\alpha}$  is the word in  $\Sigma^*$  corresponding to  $\alpha$ . Note that the set  $\varphi(r) \in \mathcal{P}(\Sigma^*)$  is prefix-closed for any  $r \in (0, 1)$ , and that if  $w \in \varphi(r)$  then also  $w\bar{0} \in \varphi(r)$  (the latter holds since, by real numbers conventions,  $0.\alpha = 0.\alpha 0$ ), so  $\varphi(r) \in \text{PersistentSafety}^*$ . Since the set of rationals with finite number of decimals in binary representation is dense in  $\mathbb{R}$  (i.e., it intersects any open interval in  $\mathbb{R}$ ) and in particular in the interval  $(0, 1)$ , it follows that the function  $\varphi : (0, 1) \rightarrow \text{PersistentSafety}^*$  is injective: indeed, if  $r_1 \neq r_2 \in (0, 1)$ , say  $r_1 < r_2$ , then there is some  $\alpha \in \{0, 1\}^*$  such that  $r_1 < 0.\alpha < r_2$ , so  $\varphi(r_1) \neq \varphi(r_2)$ . Since the interval  $(0, 1)$  has the cardinal of the continuum  $c$ , the existence of the injective function  $\varphi$  implies that  $c \leq |\text{PersistentSafety}^*|$ . By the Cantor-Bernstein-Schroeder theorem it follows that  $|\text{PersistentSafety}^*| = |\text{Safety}^*| = c$ .  $\square$

---

**From safety:** *The proof above could have been rearranged to avoid the need to use the set  $\text{PersistentSafety}^*$ . However, we prefer to keep it for two reasons:*

1. *For finite-traces, persistent safety properties appear to be more natural in the context of reactive systems than just prefix closed properties;*
  2. *Persistent safety properties play a technical bridge role in the next section to show that the infinite-trace safety properties also have the cardinal  $c$ .*
- 

With regards to finite-traces, persistent safety properties appear to be more natural in the context of reactive systems than just prefix-closed properties. Also, persistent safety properties play a technical bridge role in the next section to show that the infinite-trace safety properties also have the cardinal  $c$ .

## 3.2 Infinite Traces

The finite-trace safety properties defined above, persistent or not, rely on the intuition of a correct prefix: a safety property is identified with the set of all its finite prefixes. In the case of a persistent safety property, each “informal”

infinite acceptable behavior is captured by its infinite set of finite prefixes. Even though persistent safety properties appear to capture well in a finite-trace setting the intuition of safety in the context of (infinite-trace) reactive systems, one could argue that it does not say anything about unacceptable infinite traces. Indeed, one may think that persistent safety properties do not capture the intuition that if an infinite trace is unacceptable then there must be some finite prefix of it which is already unacceptable. In this section we show that there is in fact a bijection between safety properties over infinite traces and persistent safety properties over finite traces as we defined them in the previous section.

We start by extending the `prefixes` function to infinite traces:

**Definition 5** *Let  $\text{prefixes}: \Sigma^\omega \rightarrow \mathcal{P}(\Sigma^*)$  be the function returning for any infinite trace  $u$  all its finite prefixes  $\text{prefixes}(u)$ , and let  $\text{prefixes}: \mathcal{P}(\Sigma^\omega) \rightarrow \mathcal{P}(\Sigma^*)$  be its corresponding extension to sets of infinite traces.*

Note that  $\text{prefixes}(S) \in \text{PersistentSafety}^*$  for any  $S \in \mathcal{P}(\Sigma^\omega)$ , so `prefixes` is in fact a function  $\mathcal{P}(\Sigma^\omega) \rightarrow \text{PersistentSafety}^*$ .

The definition of safety properties over infinite traces below appears to be the most used definition of a safety property in the literature; at our knowledge, it was formally introduced by Alpern and Schneider [5], but they credit the insights of their definition to Lamport [85].

**Definition 6** *Let  $\text{Safety}^\omega$  be the set of infinite-trace properties  $Q \in \mathcal{P}(\Sigma^\omega)$  s.t.: if  $u \notin Q$  then there is a finite trace  $w \in \text{prefixes}(u)$  s.t.  $wv \notin Q$  for any  $v \in \Sigma^\omega$ .*

In other words, if an infinite behavior violates the safety property then there is some finite-trace “violation threshold”; once the violation threshold is reached, there is no chance to recover.

The following proposition can serve as an alternative and more compact definition of  $\text{Safety}^\omega$ :

**Proposition 4**  $\text{Safety}^\omega = \{Q \in \mathcal{P}(\Sigma^\omega) \mid u \in Q \text{ iff } \text{prefixes}(u) \subseteq \text{prefixes}(Q)\}.$

**Proof:** Since  $u \in Q$  implies  $\text{prefixes}(u) \subseteq \text{prefixes}(Q)$ , the only thing left to show is that  $Q \in \text{Safety}^\omega$  iff “ $\text{prefixes}(u) \subseteq \text{prefixes}(Q)$  implies  $u \in Q$ ”; the latter is equivalent to “ $u \notin Q$  implies  $\text{prefixes}(u) \not\subseteq \text{prefixes}(Q)$ ”, which is further equivalent to “ $u \notin Q$  implies there is some  $w \in \text{prefixes}(u)$  s.t.  $w \notin \text{prefixes}(Q)$ ”, which is indeed equivalent to  $Q \in \text{Safety}^\omega$ .  $\square$

Another common intuition for safety properties over infinite traces is as *closed* sets in the topology corresponding to  $\Sigma^\omega$ . Alpern and Schneider captured formally this intuition for the first time in [5]; then it was used as a convenient definition of safety by Abadi and Lamport [2, 3] among others:

**Definition 7** *An infinite sequence  $u^{(1)}, u^{(2)}, \dots$ , of infinite traces in  $\Sigma^\omega$  converges to  $u \in \Sigma^\omega$ , or  $u$  is a limit of  $u^{(1)}, u^{(2)}, \dots$ , written  $u = \lim_i u^{(i)}$ , iff for all  $m \geq 0$  there is an  $n \geq 0$  such that  $u_1^{(i)} u_2^{(i)} \dots u_m^{(i)} = u_1 u_2 \dots u_m$  for all  $i \geq n$ . If  $Q \in \mathcal{P}(\Sigma^\omega)$  then  $\overline{Q}$ , the closure of  $Q$ , is the set  $\{\lim_i u^{(i)} \mid u^{(i)} \in Q \text{ for all } i \in \mathbb{N}\}$ .*

It can be easily shown that the overline closure above is indeed a closure operator on  $\Sigma^\omega$ , that is, it is extensive ( $Q \subseteq \overline{Q}$ ), monotone ( $Q \subseteq Q'$  implies  $\overline{Q} \subseteq \overline{Q'}$ ), and idempotent ( $\overline{\overline{Q}} = \overline{Q}$ ); see Exercise 6.

**Definition 8** *Let  $\text{Safety}_{\lim}^\omega$  be the set of properties  $\{Q \in \mathcal{P}(\Sigma^\omega) \mid Q = \overline{Q}\}$ .*

As expected, the two infinite-trace safety property definitions are equivalent; we have not found any formal proof in the literature, so for the sake of completeness we give a simple proof here:

**Proposition 5**  $\text{Safety}_{\lim}^\omega = \text{Safety}^\omega$ .

**Proof:** All we need to prove is that for any  $Q \in \mathcal{P}(\Sigma^\omega)$  and any  $u \in \Sigma^\omega$ ,  $\text{prefixes}(u) \subseteq \text{prefixes}(Q)$  iff  $u = \lim_i u^{(i)}$  for some infinite sequence of infinite traces  $u^{(1)}, u^{(2)}, \dots$  in  $Q$ . If  $\text{prefixes}(u) \subseteq \text{prefixes}(Q)$  then one can find for each  $i \geq 0$  some  $u^{(i)} \in Q$  such that  $u_1 u_2 \dots u_i = u_1^{(i)} u_2^{(i)} \dots u_i^{(i)}$ , so for each  $m \geq 0$  one can pick  $n = m$  such that  $u_1 u_2 \dots u_m = u_1^{(i)} u_2^{(i)} \dots u_m^{(i)}$  for all  $i \geq n$ , so  $u = \lim_i u^{(i)}$ . Conversely, if  $u = \lim_i u^{(i)}$  for some infinite sequence of infinite traces  $u^{(1)}, u^{(2)}, \dots$  in  $Q$ , then for any  $m \geq 0$  there is some  $n \geq 0$  such that  $u_1 u_2 \dots u_m = u_1^{(n)} u_2^{(n)} \dots u_m^{(n)}$ , that is, for any prefix of  $u$  there is some  $u' \in Q$  having the same prefix, that is,  $\text{prefixes}(u) \subseteq \text{prefixes}(Q)$ .  $\square$

The next result establishes the relationship between infinite-trace safety properties and finite-trace persistent safety properties, by proposing a concrete bijective mapping relating the two (as opposed to using cardinality arguments to indirectly show only the existence of such a mapping). Therefore, there is also a bijective correspondence between safety properties over infinite traces and the real numbers:

---

*Note there are  $2^c$  properties over  $\Sigma^\omega$*

---

**Theorem 3**  $|\text{Safety}^\omega| = |\text{PersistentSafety}^*| = c$ .

**Proof:** We show that there is a bijective function between the two sets of safety properties. Recall that  $\text{prefixes}(S) \in \text{PersistentSafety}^*$  for any  $S \in \mathcal{P}(\Sigma^\omega)$ , that is, that  $\text{prefixes}$  is a function  $\mathcal{P}(\Sigma^\omega) \rightarrow \text{PersistentSafety}^*$ . Let  $\text{prefixes} : \text{Safety}^\omega \rightarrow \text{PersistentSafety}^*$  be the restriction of this prefix function to  $\text{Safety}^\omega$ . Let us also define a function  $\omega : \text{PersistentSafety}^* \rightarrow \text{Safety}^\omega$  as follows:  $\omega(P) = \{u \in \Sigma^\omega \mid \text{prefixes}(u) \subseteq P\}$ . This function is well-defined: if  $u \notin \omega(P)$  then by the definition of  $\omega(P)$  there is some  $w \in \text{prefixes}(u)$  such that  $w \notin P$ ; since  $w \in \text{prefixes}(wv)$  for any  $v \in \Sigma^\omega$ , it follows that  $wv \notin \omega(P)$  for any  $v \in \Sigma^\omega$ .

We next show that  $\text{prefixes}$  and  $\omega$  are inverse to each other. Let us first show that  $\text{prefixes}(\omega(P)) = P$  for any  $P \in \text{PersistentSafety}^*$ . The inclusion  $\text{prefixes}(\omega(P)) \subseteq P$  follows by the definition of  $\omega(P)$ :  $\text{prefixes}(u) \subseteq P$  for any  $u \in \omega(P)$ . The inclusion  $P \subseteq \text{prefixes}(\omega(P))$  follows from the fact that  $P$  is a persistent safety property: for any  $w \in P$  one can iteratively build an infinite sequence  $v_1, v_2, \dots$ , such that  $wv_1, wv_1v_2, \dots \in P$ , so  $wv_1v_2\dots \in \omega(P)$ . Let us now show that  $\omega(\text{prefixes}(Q)) = Q$  for any  $Q \in \text{Safety}^\omega$ . The inclusion  $Q \subseteq \omega(\text{prefixes}(Q))$  is immediate. For the other inclusion, let  $u \in \omega(\text{prefixes}(Q))$ , that is,  $\text{prefixes}(u) \subseteq \text{prefixes}(Q)$ . Suppose by contradiction that  $u \notin Q$ . Then there is some  $w \in \text{prefixes}(u)$  such that  $wv \notin Q$  for any  $v \in \Sigma^\omega$ . Since  $w \in \text{prefixes}(u)$  and  $\text{prefixes}(u) \subseteq \text{prefixes}(Q)$ , it follows that  $w \in \text{prefixes}(Q)$ , that is, that there is some  $u' \in Q$  such that  $u' = wv$  for some  $v \in \Sigma^\omega$ . This contradicts the fact that  $wv \notin Q$  for any  $v \in \Sigma^\omega$ . Consequently,  $u \in Q$ .

The second part follows by Theorem 2.  $\square$

### 3.3 Finite and Infinite Traces

It is also common to define safety properties as properties over both finite and infinite traces, the intuition for the finite traces being that of unfinished computations. For example, Lamport [86] extends the notion of safety in Definition 6 to properties over both finite and infinite traces, while Schneider et al [113, 54] give an alternative definition of safety over finite and infinite traces. We define both approaches shortly and then show their equivalence and their bijective correspondence with real numbers. Before that, we argue that the mix of finite and infinite traces is less trivial than it may appear,

by showing that there are significantly more prefix closed properties than in the case when only finite traces were considered.

**Definition 9** Let  $\text{PrefixClosed}^{*,\omega}$  be the set of prefix-closed sets of finite and infinite traces: for  $Q \subseteq \Sigma^* \cup \Sigma^\omega$ ,  $Q \in \text{PrefixClosed}^{*,\omega}$  iff  $\text{prefixes}(Q) \subseteq Q$ . Also, let  $\text{PersistentPrefixClosed}^{*,\omega}$  be the set of persistent prefix-closed sets of finite and infinite traces: for  $Q \in \text{PrefixClosed}^{*,\omega}$ , it is the case that  $Q \in \text{PersistentPrefixClosed}^{*,\omega} \iff$  if  $Q(w)$  for some  $w \in \Sigma^*$  then that there is some  $a \in \Sigma$  such that  $Q(wa)$ .

The next result says that there is a bijective correspondence between prefix-closed and persistent prefix-closed properties also in the case of finite and infinite traces, but that there are exponentially more such properties than in the case of just finite traces:

**Proposition 6**  $|\text{PersistentPrefixClosed}^{*,\omega}| = |\text{PrefixClosed}^{*,\omega}| = 2^c$ .

**Proof:** We show  $2^c \leq |\text{PersistentPrefixClosed}^{*,\omega}| \leq |\text{PrefixClosed}^{*,\omega}| \leq 2^c$ , where the middle inequality is immediate. For  $2^c \leq |\text{PersistentPrefixClosed}^{*,\omega}|$ , let us define  $\varphi: \mathcal{P}((0, 1)) \rightarrow \text{PersistentPrefixClosed}^{*,\omega}$  as

$$\varphi(R) = \bigcup_{0.\alpha \in R} \{\bar{\alpha}\} \cup \text{prefixes}(\bar{\alpha})$$

where we assume for any real number in the interval  $(0, 1)$  its decimal binary representation  $0.\alpha$  with  $\alpha \in \{0, 1\}^\omega$  (if the number is rational then  $\alpha$  may contain infinitely many ending 0's), and  $\bar{\alpha}$  is the infinite trace in  $\Sigma^\omega$  replacing each 0 and 1 in  $\alpha$  by  $\bar{0}$  and  $\bar{1}$ , respectively, where  $\bar{0}$  and  $\bar{1}$  are two arbitrary but fixed distinct elements in  $\Sigma$  (recall that  $|\Sigma| \geq 2$ ). Note that  $\varphi(R)$  is well-defined: it is clearly prefix-closed and it is also persistent because its finite traces are exactly prefixes of infinite traces, so they admit continuations in  $\varphi(R)$ . It is easy to see that  $\varphi$  is injective. Since  $|(0, 1)| = c$ , we conclude that  $2^c \leq |\text{PersistentPrefixClosed}^{*,\omega}|$ .

To show  $|\text{PrefixClosed}^{*,\omega}| \leq 2^c$ , note that any property in  $\text{PrefixClosed}^{*,\omega}$  is a union of a subset in  $\Sigma^*$  and a subset in  $\Sigma^\omega$ , so  $|\text{PrefixClosed}^{*,\omega}| \leq 2^{|\Sigma^*|} \cdot 2^{|\Sigma^\omega|}$ . Since  $|\Sigma^*| = \aleph_0$ ,  $|\Sigma^\omega| = c$ ,  $2^{\aleph_0} = c$ , and  $c \cdot 2^c = 2^c$  (by absorption of transfinite cardinals), we get that  $|\text{PrefixClosed}^{*,\omega}| \leq 2^c$ .  $\square$

The fact that properties in  $\text{PersistentPrefixClosed}^{*,\omega}$  contain also infinite traces was crucial in showing the injectivity of  $\varphi$  in the proof above. A similar construction for the finite trace setting does *not* work. Indeed, if

one tries to define a function  $\varphi: \mathcal{P}((0, 1)) \rightarrow \text{PersistentSafety}^*$  as  $\varphi(R) = \bigcup_{0.\alpha \in R} \text{prefixes}(\bar{\alpha})$ , then one can show it well-defined but cannot show it injective: e.g.,  $\varphi((0, 0.5)) = \varphi((0, 0.5])$ .

Since safety properties over finite and infinite traces are governed by the same intuitions as safety properties over only finite or over only infinite traces, the result above tells us that prefix closeness is not a sufficient condition to properly capture the safety properties. Schneider [113] proposes an additional condition in the context of his EM (execution monitoring) framework, namely that if an infinite trace is not in the property, then there is a finite prefix of it which is not in the property either. It is easy to see that this additional condition is equivalent to saying that an infinite trace is in the property whenever all its finite prefixes are in the property, which allows us to compactly define safety properties over finite and infinite traces in the EM style as follows:

**Definition 10**  $\text{Safety}_{\text{EM}}^{*,\omega} = \{Q \subseteq \Sigma^* \cup \Sigma^\omega \mid u \in Q \text{ iff } \text{prefixes}(u) \subseteq Q\}$ .

Note that  $\text{Safety}_{\text{EM}}^{*,\omega} \subset \text{PrefixClosed}^{*,\omega}$ . We will shortly show that  $\text{Safety}_{\text{EM}}^{*,\omega}$  is in fact much smaller than  $\text{PrefixClosed}^{*,\omega}$ , by showing that  $|\text{Safety}_{\text{EM}}^{*,\omega}| = c$ .

The consecrated definition of a safety property in the context of both finite and infinite traces is perhaps the one proposed by Lamport in [86], which relaxes the one in Definition 6 by allowing  $u$  to range over both finite and infinite traces:

**Definition 11** *Let  $\text{Safety}^{*,\omega}$  be the set of finite- and infinite-trace properties  $\{Q \subseteq \Sigma^* \cup \Sigma^\omega \mid u \notin Q \Rightarrow (\exists w \in \text{prefixes}(u)) (\forall v \in \Sigma^* \cup \Sigma^\omega) wv \notin Q\}$*

Schneider informally stated in [113] that the two definitions of safety above are equivalent. It is not hard to show it formally:

**Proposition 7**  $\text{Safety}_{\text{EM}}^{*,\omega} = \text{Safety}^{*,\omega}$ .

**Proof:** First note that  $\text{Safety}^{*,\omega} \subseteq \text{PrefixClosed}^{*,\omega}$ : if  $wu \in Q \in \text{Safety}^{*,\omega}$  and  $w \notin Q$  then there is some  $w' \in \text{prefixes}(w)$ , say  $w = w'w''$ , such that  $w'v \notin Q$  for any  $v$ , in particular  $w'w''u \notin Q$ , which contradicts  $wu \in Q$ .

$\text{Safety}^{*,\omega} \subseteq \text{Safety}_{\text{EM}}^{*,\omega}$ : let  $Q \in \text{Safety}^{*,\omega}$  and  $u \in \Sigma^* \cup \Sigma^\omega$  s.t.  $\text{prefixes}(u) \subseteq Q$ ; if  $u \notin Q$  then there is some  $w \in \text{prefixes}(u)$  s.t.  $wv \notin Q$  for any  $v$ , in particular for  $v$  the empty word, that is,  $w \notin Q$ , which contradicts  $\text{prefixes}(u) \subseteq Q$ .

$\text{Safety}_{\text{EM}}^{*,\omega} \subseteq \text{Safety}^{*,\omega}$ : let  $u \notin Q \in \text{Safety}_{\text{EM}}^{*,\omega}$ ; then  $\text{prefixes}(u) \not\subseteq Q$ , that is, there is some  $w \in \text{prefixes}(u)$  s.t.  $w \notin Q$ ; since  $Q$  is prefix-closed, it follows that  $wv \notin Q$  for any  $v \in \Sigma^* \cup \Sigma^\omega$ .  $\square$

We next show that there is a bijective correspondence between the safety properties over finite or infinite traces above and the finite trace safety properties in Section 3.1:

**Theorem 4**  $|\text{Safety}^{*,\omega}| = |\text{Safety}_{\text{EM}}^{*,\omega}| = |\text{Safety}^*| = c$ .

**Proof:**  $\text{Safety}^* \subset \text{Safety}_{\text{EM}}^{*,\omega}$  since the properties in  $\text{Safety}_{\text{EM}}^{*,\omega}$  are prefix-closed, so  $|\text{Safety}^*| \leq |\text{Safety}_{\text{EM}}^{*,\omega}|$ .

Since the functions  $\text{prefixes}: \mathcal{P}(\Sigma^*) \rightarrow \mathcal{P}(\Sigma^*)$  and  $\text{prefixes}: \mathcal{P}(\Sigma^\omega) \rightarrow \mathcal{P}(\Sigma^*)$  have actual co-domains  $\text{Safety}^*$  and  $\text{PersistentSafety}^*$ , respectively, they can be organized as a function  $\text{prefixes}: \text{Safety}_{\text{EM}}^{*,\omega} \rightarrow \text{Safety}^*$ . Let us show that this function is injective. Let us assume  $Q \neq Q' \in \text{Safety}_{\text{EM}}^{*,\omega}$ , say  $u \in Q$  and  $u \notin Q'$ , s.t.  $\text{prefixes}(Q) = \text{prefixes}(Q')$ . Since  $u \in Q \in \text{Safety}_{\text{EM}}^{*,\omega}$  it follows that  $\text{prefixes}(u) \subseteq \text{prefixes}(Q) \subseteq Q$ , which implies that  $\text{prefixes}(u) \subseteq \text{prefixes}(Q') \subseteq Q'$ ; since  $Q' \in \text{Safety}_{\text{EM}}^{*,\omega}$ , it follows that  $u \in Q'$ , contradiction. Therefore,  $\text{prefixes}: \text{Safety}_{\text{EM}}^{*,\omega} \rightarrow \text{Safety}^*$  is injective, which proves that  $\text{Safety}_{\text{EM}}^{*,\omega} \leq \text{Safety}^*$ .

The rest follows by Proposition 7 and Theorem 2.  $\square$

### 3.4 “Always Past” Characterization

Another common way to specify safety properties is by giving an arbitrary property on finite traces, not necessarily prefix closed, and then to require that any acceptable behavior must have all its finite prefixes in the given property. A particularly frequent case is when one specifies the property of the finite-prefixes using the past-time fragment of linear temporal logics (LTL). For example, Manna and Pnueli [89] call the resulting “always (past LTL)” properties *safety formulae*; many other authors, including ourselves, adopted the terminology “safety formula” from Manna and Pnueli, although some qualify it as “LTL safety formula”. An example of an LTL safety formula is “always ( $b$  implies eventually in the past  $a$ )”, written using LTL notation as “ $\square(b \rightarrow \diamond a)$ ”; here the past time formula “ $b \rightarrow \diamond a$ ” compactly specifies all the finite-traces

$$\{ws w' s' \mid w, w' \in \Sigma^*, s, s' \in \Sigma, a(s) \text{ and } b(s') \text{ hold}\} \cup \{ws \mid w \in \Sigma^*, s \in \Sigma, b(s) \text{ does not hold}\}.$$

---

**From safety:** *We will investigate the case when safety properties are expressed as LTL safety formulae, as well as optimal monitoring techniques for such safety properties, in Section 10.1.*

---

In the remainder of this section we assume that the past time prefix properties are given as ordinary sets of finite-traces (so we make abstraction of how these properties are expressed) and show not only that the resulting “always past” properties are safety properties, but also that any safety properties can be expressed as an “always past” property. This holds for all the variants of safety properties (i.e., over finite traces, over infinite traces, or over both finite and infinite traces).

**Definition 12** *Let  $P \subseteq \Sigma^*$  be any property over finite traces. Then we define the “always past” property  $\Box P$  as follows:*

- (finite traces)  $\{w \in \Sigma^* \mid \text{prefixes}(w) \subseteq P\}$ ; and
- (infinite traces)  $\{u \in \Sigma^\omega \mid \text{prefixes}(u) \subseteq P\}$ ; and
- (finite and infinite traces)  $\{u \in \Sigma^* \cup \Sigma^\omega \mid \text{prefixes}(u) \subseteq P\}$ .

*Let  $\text{Safety}_\Box^*$ ,  $\text{Safety}_\Box^\omega$  and  $\text{Safety}_\Box^{*,\omega}$  be the corresponding sets of properties.*

---

**From safety:** *In Section 10.1 we show that the language  $\mathcal{L}(\Box\varphi)$ , that corresponds to the LTL safety formula  $\Box\varphi$  for  $\varphi$  some past-time LTL formula, is a property in  $\text{Safety}_\Box^\omega$ . If one was interested in a finite-trace or in a both finite and infinite trace semantics of LTL, then one could have shown that  $\mathcal{L}(\Box\varphi) \in \text{Safety}_\Box^*$  or that  $\mathcal{L}(\Box\varphi) \in \text{Safety}_\Box^{*,\omega}$ .*

---

Intuitively, one can regard the square “ $\Box$ ” as a closure operator. Technically, it is not precisely a closure operator because it does not operate on the same set: it takes finite-trace properties to any of the three types of properties considered. Since  $\text{prefixes}$  takes properties back to finite-trace properties, we can show the following result saying that the square is a “closure operator via  $\text{prefixes}$ ”, and that safety properties are precisely the sets of words which are closed this way:

**Proposition 8** *The following hold for all three types of safety properties:*

- $\Box(\text{prefixes}(\Box P)) = \Box P$  for any  $P \subseteq \Sigma^*$ ;
- $Q$  is a safety property iff  $\Box(\text{prefixes}(Q)) = Q$ .

**Proof:** Left as an exercise to the reader. See Exercise 5.  $\square$

We next show that the “always past” properties are all safety properties and, moreover, that any safety property can be expressed as an “always past” property:

**Theorem 5** *The following hold:*

- $\text{Safety}_{\square}^* = \text{Safety}^*$ ,
- $\text{Safety}_{\square}^{\omega} = \text{Safety}^{\omega}$ , and
- $\text{Safety}_{\square}^{*,\omega} = \text{Safety}^{*,\omega}$ .

Therefore, each of the “always past” safety properties have the cardinal  $c$ .

**Proof:** We prove each of the equalities by double inclusion.

$\text{Safety}_{\square}^* \subseteq \text{Safety}^*$ . It is true because any property  $\square P$  in  $\text{Safety}_{\square}^*$  is prefix-closed.

$\text{Safety}^* \subseteq \text{Safety}_{\square}^*$ . If  $P \in \text{Safety}^*$  then we claim that  $P = \square P$ , so  $P \in \text{Safety}_{\square}^*$ . Indeed, since  $P$  is prefix-closed,  $\text{prefixes}(w) \subseteq P$  for any  $w \in P$ , so  $w \in \square P$ ; also, since  $w \in \text{prefixes}(w)$ , it follows that for any  $w \in \square P$ ,  $w \in P$ .

$\text{Safety}_{\square}^{\omega} \subseteq \text{Safety}^{\omega}$ . Let  $\square P$  be an “always past” property in  $\text{Safety}_{\square}^{\omega}$ , and let  $u$  be an infinite trace in  $\Sigma^{\omega}$  such that  $u \notin \square P$ . Then it follows that  $\text{prefixes}(u) \not\subseteq P$ , that is, there is some  $w \in \text{prefixes}(u)$  such that  $w \notin P$ . Since  $w \in \text{prefixes}(wv)$  for any  $v \in \Sigma^{\omega}$ , it means that there is no  $v \in \square P$  such that  $\text{prefixes}(wv) \subseteq P$ , that is, there is no  $v \in \Sigma^{\omega}$  such that  $wv \in \square P$ . Therefore,  $\square P \in \text{Safety}^{\omega}$ .

$\text{Safety}^{\omega} \subseteq \text{Safety}_{\square}^{\omega}$ . If  $Q \in \text{Safety}^{\omega}$  then we claim that  $Q = \square \text{prefixes}(Q)$ . The inclusion  $Q \subseteq \square \text{prefixes}(Q)$  is clear, because  $u \in Q$  implies  $\text{prefixes}(u) \subseteq \text{prefixes}(Q)$ . For the other inclusion, note that if  $\text{prefixes}(u) \subseteq \text{prefixes}(Q)$  for some  $u \in \Sigma^{\omega}$ , then  $u$  must be in  $Q$ : if  $u \notin Q$  then by the definition of  $Q \in \text{Safety}^{\omega}$ , there is some  $w \in \text{prefixes}(u)$  which cannot be completed into an infinite trace in  $Q$ , which contradicts  $\text{prefixes}(u) \subseteq \text{prefixes}(Q)$ .

$\text{Safety}_{\square}^{*,\omega} \subseteq \text{Safety}^{*,\omega}$ . By Proposition 7, it suffices to show that  $\text{Safety}_{\square}^{*,\omega} \subseteq \text{Safety}_{\text{EM}}^{*,\omega}$ . Let  $\square P$  be an “always past” property in  $\text{Safety}_{\square}^{*,\omega}$ , and let  $u \in \Sigma^* \cup \Sigma^{\omega}$  such that  $\text{prefixes}(u) \subseteq \text{prefixes}(\square P)$ . Since  $\text{prefixes}(\square P) \subseteq P$ , it follows that  $u \in \square P$ ; therefore,  $\square P \in \text{Safety}_{\text{EM}}^{*,\omega}$ .

$\text{Safety}^{*,\omega} \subseteq \text{Safety}_{\square}^{*,\omega}$ . It is straightforward to see that  $Q \in \text{Safety}_{\text{EM}}^{*,\omega}$  implies  $Q = \square \text{prefixes}(Q)$ .

The cardinality part follows by Theorems 2, 3, and 4.  $\square$

Proposition 8 and Theorem 5 give yet another characterization for safety properties over any of the three combinations of traces, namely one in the style of the equivalent formulation of safety over infinite traces in Proposition 4:  $Q$  is a safety property iff it contains precisely the words whose prefixes are in  $\text{prefixes}(Q)$ .

## Exercises

**Exercise 2** The  $\text{prefixes}: \mathcal{P}(\Sigma^*) \rightarrow \mathcal{P}(\Sigma^*)$  in Definition 1 is a closure operator: it is extensive ( $P \subseteq \text{prefixes}(P)$ ), monotone ( $P \subseteq P'$  implies  $\text{prefixes}(P) \subseteq \text{prefixes}(P')$ ), and idempotent ( $\text{prefixes}(\text{prefixes}(P)) = \text{prefixes}(P)$ ).

**Exercise 3** (Counter-)Example 4 showed that the finiteness of  $\Sigma$  was necessary in order for Proposition 2 to hold, by defining a property  $P$  over  $\Sigma = \mathbb{N} \cup \{\infty\}$  in which all non-empty words start with  $\infty$ . Can we remove  $\infty$  from  $\Sigma$  and from all the words in  $P$ ? Why, or why not?

**Exercise 4** Same like Exercise 3, but for Example 5 instead of Example 4.

**Exercise 5** Prove Proposition 8.

**Exercise 6** The “closure under limits” operation in Definition 7 is indeed a closure operator on  $\Sigma^\omega$ : it is extensive ( $Q \subseteq \overline{Q}$ ), monotone ( $Q \subseteq Q'$  implies  $\overline{Q} \subseteq \overline{Q'}$ ), and idempotent ( $\overline{\overline{Q}} = \overline{Q}$ ).

## Chapter 4

# Monitoring

In this section we give yet another characterization of safety properties, namely as monitorable properties. Specifically, we formally define a monitor as a (possibly infinite) state machine without final states but with a partial transition function, and then we show that safety properties are precisely the properties that can be monitored with such monitors. We then elaborate on the problem of defining the complexity of monitoring a safety property, discussing some pitfalls and guiding principles, and show that monitoring a safety property can be an arbitrarily hard problem. Finally, we give a more compact and mathematical equivalent definition of a monitor, which may be useful in further foundational efforts in this area.

---

*Relate our definition of a monitor with Schneider's security automata*

---

### 4.1 Specifying Safety Properties as Monitors

Safety properties are difficult to work with as flat sets of finite or infinite words, not only because they can contain infinitely many words, but also because such a flat representation is inconvenient for further analysis. It is important therefore to *specify* safety properties using formalisms that are easier to represent and reason about.

---

**From safety:** *The next sections in this paper investigate several dedicated formalisms that proved to be convenient in specifying safety, such as finite state machines, regular expressions and temporal logics, together with corresponding limitations and efficient monitor synthesis techniques.*

---

Formalisms known to be useful for specifying safety properties include regular expressions and temporal logics, which can be efficiently translated into finite-state machines which can then be used as monitors. In this section we formalize the intuitive notion of a *monitor* as a special state machine and give yet another characterization of safety properties, namely as *monitorable properties*. Since monitorable properties are completely defined by their monitors, it follows that *all* safety properties can be specified by their corresponding monitors.

Recall that we work under the assumption that  $\Sigma$  is a set of events or program states such that  $|\Sigma| \leq \aleph_0$ .

**Definition 13** *A  $\Sigma$ -monitor, or just a monitor (when  $\Sigma$  is understood), is a triple  $\mathcal{M} = (S, s_0, M : S \times \Sigma \rightarrow S)$ , where  $S$  is a set of states,  $s_0 \in S$  is the initial state, and  $M$  is a deterministic partial transition function.*

Therefore, a monitor as defined above is nothing but a deterministic state machine without final states. Moreover, the set of states is allowed to be infinite, and the transition function has no complexity requirements (it can even be undecidable). We could have defined monitors to be standard state machines, but the subsequent technical developments would have been slightly more involved.

---

**From safety:** *In fact, we aim at shortly giving an even more compact definition of a monitor, that we will call canonical monitor, which appears to be sufficient to capture any safety property.*

---

The intuition for a monitor is the expected one: the monitor is driven by events generated by the observed program (the letters in  $\Sigma$ )—each newly received event drives the monitor from its current state to some other state, as indicated by the transition function  $M$ ; if the monitor ever gets stuck, that is, if the transition function  $M$  is undefined on the current state and the current event, then the monitored property is declared violated at that point by the monitor.

For any partial function  $M: S \times \Sigma \rightarrow S$ , we obey the following common notational convention. If  $s \in S$  and  $w = w_1 w_2 \dots w_k \in \Sigma^*$ , we write “ $M(s, w) \downarrow$ ” whenever  $M(s, w)$  is defined, that is, whenever  $M(s, w_1)$  and  $M(M(s, w_1), w_2)$  and ... and  $M(\dots(M(s, w_1), w_2) \dots, w_k)$  are all defined, which is nothing but only saying that  $M(\dots(M(s, w_1), w_2) \dots, w_k)$  is defined. If we write  $M(s, w) = s'$  for some  $s' \in S$ , then, as expected, we mean that  $M(\dots(M(s, w_1), w_2) \dots, w_k)$  is defined and equal to  $s'$ .

A monitor specifies a finite-trace property, an infinite-trace property, as well as a finite- and infinite-trace property:

**Definition 14** *Given a monitor  $\mathcal{M} = (S, s_0, M: S \times \Sigma \rightarrow S)$ , we define the following properties:*

- $\mathcal{L}^*(\mathcal{M}) = \{w \in \Sigma^* \mid M(s_0, w) \downarrow\}$ ,
- $\mathcal{L}^\omega(\mathcal{M}) = \{u \in \Sigma^\omega \mid M(s_0, w) \downarrow \text{ for all } w \in \text{prefixes}(u)\}$ , and
- $\mathcal{L}^{*,\omega}(\mathcal{M}) = \mathcal{L}^*(\mathcal{M}) \cup \mathcal{L}^\omega(\mathcal{M})$ .

We call  $\mathcal{L}^*(\mathcal{M})$  the finite-trace property specified by  $\mathcal{M}$ , call  $\mathcal{L}^\omega(\mathcal{M})$  the infinite-trace property specified by  $\mathcal{M}$ , and call  $\mathcal{L}^{*,\omega}(\mathcal{M})$  the finite- and infinite-trace property specified by  $\mathcal{M}$ . Also, we let

$$\mathcal{S}_{\mathcal{M}} = \{s \in S \mid (\exists w \in \Sigma^*) M(s_0, w) = s\}$$

denote the set of reachable states of  $\mathcal{M}$ .

A *monitorable* property is a property which can be specified by a monitor. We next capture this intuitive notion formally:

**Definition 15** *For a property  $P \subseteq \Sigma^* \cup \Sigma^\omega$ , we let  $\text{Monitors}(P)$  be the set of monitors  $\{\mathcal{M} \mid \mathcal{L}^{*,\omega}(\mathcal{M}) = P\}$ . If  $\text{Monitors}(P) \neq \emptyset$  then  $P$  is called monitorable and the elements of  $\text{Monitors}(P)$  are called monitors of  $P$ . We define the following classes of properties:*

- $\text{Monitorable}^* = \{P \subseteq \Sigma^* \mid P \text{ monitorable}\}$ ,
- $\text{Monitorable}^\omega = \{P \subseteq \Sigma^\omega \mid P \text{ monitorable}\}$ , and
- $\text{Monitorable}^{*,\omega} = \{P \subseteq \Sigma^* \cup \Sigma^\omega \mid P \text{ monitorable}\}$ .

The notion of persistence can also be adapted to monitors:

**Definition 16** A monitor  $\mathcal{M} = (S, s_0, M : S \times \Sigma \rightarrow S)$  is persistent iff for any reachable state  $s \in \mathcal{S}_{\mathcal{M}}$ , there is an  $a \in \Sigma$  such that  $M(s, a) \downarrow$ . Let

- $\text{PersistentMonitorable}^* = \{\mathcal{L}^*(\mathcal{M}) \mid \mathcal{M} \text{ persistent}\}$

be the set of finite-trace properties monitorable by persistent monitors.

Our next goal is to show that each monitor admits a largest persistent “submonitor”. To formalize it, we lift the conventional partial order relation on partial functions to monitors:

**Definition 17** If  $\mathcal{M}_1 = (S, s_0, M_1 : S \times \Sigma \rightarrow S)$  and  $\mathcal{M}_2 = (S, s_0, M_2 : S \times \Sigma \rightarrow S)$  are two monitors sharing the same states and initial state, then let  $\mathcal{M}_1 \sqsubseteq \mathcal{M}_2$ , read  $\mathcal{M}_1$  a submonitor of  $\mathcal{M}_2$ , iff for any  $s \in S$  and any  $a \in \Sigma$ , if  $M_1(s, a)$  is defined then  $M_2(s, a)$  is also defined and  $M_2(s, a) = M_1(s, a)$ .

The above can be easily generalized to allow  $\mathcal{M}_1$  to only have a subset of the states of  $\mathcal{M}_2$ , but we found that generalization unnecessary so far.

The above partial-order on monitors allows us to use conventional mathematics to obtain the largest persistent sub-monitor of a monitor:

**Proposition 9**  $(\{\mathcal{K} \mid \mathcal{K} \sqsubseteq \mathcal{M} \text{ and } \mathcal{K} \text{ persistent}\}, \sqsubseteq)$  is a complete (join) semilattice for any monitor  $\mathcal{M}$ .

**Proof:** If  $\{\mathcal{K}_i = (S, s_0, K_i : S \times \Sigma \rightarrow S) \in \mathcal{M}\}_{i \in I}$  is a set of persistent monitors, then their supremum (or join) is the monitor  $\mathcal{K} = (S, s_0, K : S \times \Sigma \rightarrow S)$  where  $K(s, a) = s'$  iff there is some  $i \in I$  such that  $K_i(s, a) = s'$ . It is easy to see that  $\mathcal{K}$  is a well-defined monitor and that it is persistent.  $\square$

Since complete semilattices have maximum elements, the following definition is fully justified:

**Definition 18** For any monitor  $\mathcal{M} = (S, s_0, M : S \times \Sigma \rightarrow S)$ , we let  $\mathcal{M}^\circ = (S, s_0, M^\circ : S \times \Sigma \rightarrow S)$  be the  $\sqsubseteq$ -maximal element of the complete lattice  $(\{\mathcal{K} \mid \mathcal{K} \sqsubseteq \mathcal{M} \text{ and } \mathcal{K} \text{ persistent}\}, \sqsubseteq)$ .

We next show that, as expected, there is a tight relationship between persistent safety properties (Definition 3) and persistent canonical monitors.

**Proposition 10** Let  $\mathcal{M} = (S, s_0, M : S \times \Sigma \rightarrow S)$ . Then the following hold:

- $\mathcal{L}^\omega(\mathcal{M}) = \mathcal{L}^\omega(\mathcal{M}^\circ)$ ,

- $\mathcal{L}^*(\mathcal{M}^\circ) = \mathcal{L}^*(\mathcal{M})^\circ$ , and
- $\mathcal{M}$  persistent iff  $\mathcal{L}^*(\mathcal{M})$  persistent.

**Proof:** The first property can be shown by the following sequence of equivalences:  $u \in \mathcal{L}^\omega(\mathcal{M})$  iff  $M(s_0, w) \downarrow$  for all  $w \in \text{prefixes}(u)$ , iff there is some persistent monitor  $\mathcal{K} \sqsubseteq \mathcal{M}$  such as  $K(s_0, w) \downarrow$  for all  $w \in \text{prefixes}(u)$ , iff  $M^\circ(s_0, w) \downarrow$  for all  $w \in \text{prefixes}(u)$ , iff  $u \in \mathcal{L}^\omega(\mathcal{M}^\circ)$ .

The second property can be shown as follows:  $w \in \mathcal{L}^*(\mathcal{M}^\circ)$  iff  $M^\circ(s_0, w) \downarrow$ , iff there is some  $u \in \mathcal{L}^\omega(\mathcal{M}^\circ)$  such that  $w \in \text{prefixes}(u)$  (because  $\mathcal{M}^\circ$  is persistent), iff there is some  $u \in \mathcal{L}^\omega(\mathcal{M})$  such that  $w \in \text{prefixes}(u)$  (by the first property), iff there is some  $u \in \Sigma^\omega$  such that  $w \in \text{prefixes}(u) \subseteq \mathcal{L}^*(\mathcal{M})$ , iff there is some  $u \in \Sigma^\omega$  such that  $w \in \text{prefixes}(u) \subseteq \mathcal{L}^*(\mathcal{M})^\circ$  (because  $\text{prefixes}(u)$  is a persistent safety property), iff  $w \in \mathcal{L}^*(\mathcal{M})^\circ$ .

Finally, the third property is an immediate consequence of the second, noticing that  $\mathcal{M}$  is persistent iff it is equal to  $\mathcal{M}^\circ$ , and that  $\mathcal{L}^*(\mathcal{M})$  is persistent iff it is equal to  $\mathcal{L}^*(\mathcal{M})^\circ$ .  $\square$

**Theorem 6** *The following hold:*

- $\text{Monitorable}^* = \text{Safety}^*$ ,
- $\text{Monitorable}^\omega = \text{Safety}^\omega$ ,
- $\text{Monitorable}^{*,\omega} = \text{Safety}^{*,\omega}$ , and
- $\text{PersistentMonitorable}^* = \text{PersistentSafety}^*$ .

**Proof:** First, note that the following hold for any monitor  $\mathcal{M}$ :

- $\mathcal{L}^*(\mathcal{M}) \in \text{Safety}^*$ ,
- $\mathcal{L}^\omega(\mathcal{M}) \in \text{Safety}^\omega$ , and
- $\mathcal{L}^{*,\omega}(\mathcal{M}) \in \text{Safety}^{*,\omega}$ .

These all follow by Theorem 5: taking  $P$  in Definition 12 to be the property  $\{w \in \Sigma^* \mid M(s_0, w) \downarrow\}$ , then  $\Box P$  over finite traces is precisely  $\mathcal{L}^*(\mathcal{M})$ , over infinite traces is precisely  $\mathcal{L}^\omega(\mathcal{M})$ , and over finite and infinite traces is precisely  $\mathcal{L}^{*,\omega}(\mathcal{M})$ , so the three languages are in  $\text{Safety}_\Box^*$ ,  $\text{Safety}_\Box^\omega$ , and  $\text{Safety}_\Box^{*,\omega}$ , respectively. Therefore,  $\text{Monitorable}^* \subseteq \text{Safety}^*$ ,  $\text{Monitorable}^\omega \subseteq \text{Safety}^\omega$ , and  $\text{Monitorable}^{*,\omega} \subseteq \text{Safety}^{*,\omega}$ .

Second, note that we can associate a default monitor  $\mathcal{M}_P$  to any finite-trace property  $P \subseteq \Sigma^*$ , namely  $(S_P, \epsilon, M_P : S_P \times \Sigma \rightarrow S_P)$ , where  $S_P = \text{prefixes}(P)$ ,  $\epsilon$  is the empty word, and  $M_P(w, a)$  is defined iff  $wa \in \text{prefixes}(P)$ , and in that case  $M_P(w, a) = wa$ . Moreover, it is easy to check that

- $\mathcal{L}^*(\mathcal{M}_P) = \{w \in \Sigma^* \mid \text{prefixes}(w) \subseteq P\} = \Box P$  (over finite traces) ,
- $\mathcal{L}^\omega(\mathcal{M}_P) = \{u \in \Sigma^\omega \mid \text{prefixes}(u) \subseteq P\} = \Box P$  (over infinite traces),
- $\mathcal{L}^{*,\omega}(\mathcal{M}_P) = \{u \in \Sigma^* \cup \Sigma^\omega \mid \text{prefixes}(u) \subseteq P\} = \Box P$  (over both finite and infinite traces).

Since  $P$  was chosen arbitrarily, it follows then by Theorem 5 that  $\text{Safety}^* \subseteq \text{Monitorable}^*$ ,  $\text{Safety}^\omega \subseteq \text{Monitorable}^\omega$ , and  $\text{Safety}^{*,\omega} \subseteq \text{Monitorable}^{*,\omega}$ .

Finally, the equality  $\text{PersistentMonitorable}^* = \text{PersistentSafety}^*$  follows by the first fact and by Proposition 10.  $\square$

## 4.2 Complexity of Monitoring a Safety Property

We here address the problem of defining the complexity of monitoring. Before we give our definition, let us first discuss some pitfalls in defining this notion. Our definition for the complexity of monitoring resulted as a consequence of trying to avoid these pitfalls. Let  $P$  be a safety property.

### Pitfall 1.

*The complexity of monitoring  $P$  is nothing but the complexity of checking, for an input word  $w \in \Sigma^*$ , whether  $w \in \text{prefixes}(P)$ .*

This would be an easy to formulate decision problem, but, unfortunately, does not capture well the intuition of monitoring, because it does not require that the word  $w$  be processed incrementally, as its letters become available from the observed system. Incremental processing of letters can make a huge difference in both how complex monitoring is and how monitoring complexity can be defined. For example, it is well-known that the membership problem of a finite word to the language of an extended regular expression (ERE), i.e., a regular expression extended with complement operators, is a polynomial problem (the classic algorithm by Hopcroft and Ullman [70] runs in space  $O(m^2 \cdot n)$  and time  $O(m^3 \cdot n)$ , where  $m$  is the size of the word and  $n$  that of the expression). However,

---

**From safety:** as shown in Section 7

---

there are EREs defining safety properties whose monitoring requires non-elementary space and time. Of course, this non-elementary lower-bound is expressed only as a function of the size of the ERE representing the safety property; it does not take into account the size of the monitored trace. This leads us to our first guiding principle:

**Principle 1.**

The complexity of monitoring a safety property  $P$  should depend only upon  $P$ , not upon the trace being monitored.

Indeed, since monitoring is a process that involves potentially unbounded traces, if the complexity of monitoring a property  $P$  were expressed as a function of the execution trace as well, then that complexity measure would be close to meaningless in practice, because monitoring reactive systems would have unbounded complexity. For example, consider an operating system monitoring some safety property on how its resources are being used by the various running processes; what one would like to know here is what is the runtime overhead of monitoring that safety property at each relevant event, and not the obvious fact that the more the operating system runs the larger the total runtime overhead is.

Nevertheless, one can admittedly argue that it would still be useful to know how complex the monitoring of  $P$  against a given finite trace  $w$  is, in terms of both the size of (some representation of)  $P$  and the size of  $w$ ; however, this is nothing but a conventional membership test decision problem, that has nothing to do with monitoring. If one picks some arbitrary off-the-shelf efficient algorithm for membership testing and uses that at each newly received event on the existing finite execution trace, then one may obtain a “monitoring” algorithm whose complexity to process each event grows in time, as events are processed. In the context of monitoring a reactive system, that means that eventually the monitoring process may become unfeasible, regardless of how many resources are initially available and regardless of how efficient the membership testing algorithm is. What one needs in order for the monitoring process to stay feasible regardless of how many events are observed, is a special membership algorithm that processes each event as received and whose state or processing time does not increase potentially unbounded as events are received. Therefore, one needs an algorithm which, if it takes resources  $R$  to check  $w$ , then it takes at most  $R + \Delta$  to check a one-event continuation  $wa$  of  $w$ , where  $\Delta$  *does not depend on  $w$* . In other words, one needs a *monitor for  $P$  of complexity  $\Delta$* .

**Pitfall 2.**

*$P$  is typically infinite, so the complexity of monitoring  $P$  should be a function of the size of some finite specification, or representation, of  $P$ .*

Indeed, since Principle 1 tells us that the complexity of monitoring  $P$  is a function of  $P$  only and not of the monitored trace, one may be tempted to conclude that it is a function of the *size* of some convenient encoding of  $P$ . There are at least two problems with this approach, that we discuss below.

- One problem is that the same property  $P$  can be specified in many different ways as a structure of finite size; for example, it can be specified as a regular expression, as an extended regular expression, as a temporal logic formula, as an ordinary automaton, as a push-down automaton, etc. These formalisms may represent  $P$  as specifications of quite different sizes. Which is the most appropriate? It is, nevertheless, interesting and important to study the complexity of monitoring safety properties expressed using different specification formalisms, as a function of the property representation size, because that can give us an idea of the amount of resources needed to monitor a particular specification. However, one should be aware that such a complexity measure is an attribute of the corresponding specification formalisms, not of the specified property itself. Indeed, the higher this complexity measure for a particular formalism, the higher the encoding strength of safety properties in that formalism: for example, the complexity of monitoring safety properties expressed as EREs is non-elementary in the size of the original ERE, while the complexity of monitoring the same property expressed as an ordinary regular expression is linear in the size of the regular expression. Does that mean that one can monitor safety properties expressed as regular expressions non-elementarily more efficiently than one can monitor safety properties expressed as EREs? Of course not, because EREs and regular expressions have the same expressiveness, so they specify exactly the same safety properties. All it means is that EREs can express safety-properties non-elementarily more compactly than ordinary regular expressions.
- Another problem with this approach is that apparently appropriate representations of  $P$  may be significantly larger than it takes to monitor  $P$ . One may say, for example, that, whenever possible, a natural way to specify a particular safety property is as a finite-state machine, e.g.,

as a monitor like in Definition 13 . To be more concrete, consider that the safety property  $P_n$  saying “every  $2^n$ -th event is  $a$ ” is specified as a monitor of  $2^n$  states that transits with any event from each state to the next one, except for the  $2^n$ -th state, which has only one transition, with event  $a$ , back to state 1. Therefore, the size of this representation of  $P_n$  is  $\Omega(2^n)$ . Assuming that each state takes  $n$  bits of storage (for example, assume that states are exactly the binary encodings of the numbers  $1, 2, 3, \dots, 2^n$ ) and that the next state can be calculated from the current state in linear complexity with the size of the state (which is true in our scenario), then it is clear that the actual complexity of monitoring  $P_n$  is  $O(n)$ . If the complexity of monitoring  $P_n$  were a function of the size of the specification of  $P_n$ , then one could wrongly conclude that the complexity of monitoring “every  $2^n$ -th event is  $a$ ” is  $O(2^n)$ .

Therefore, a safety property  $P$  has an inherent complexity w.r.t. monitoring, complexity which has nothing to do with how  $P$  is represented, or encoded, or specified. It is that inherent complexity attribute of safety properties that we are after here. From the discussion above, we draw our second guiding principle:

**Principle 2.**

The monitoring complexity of a safety property  $P$  is an attribute of  $P$  alone, not a function of the size of some adhoc representation of  $P$ .

By Theorem 6, safety properties are precisely those properties that are monitorable, that is, those properties  $P$  for which there are (finite-state or not) monitors  $\mathcal{M} = (S, s_0, M : S \times \Sigma \rightarrow S)$  whose (finite-trace, infinite-trace, or finite- and infinite-trace—this depends upon the type of  $P$ ) language is precisely  $P$ . Any algorithm, program or system that one may come up with to be used as a monitor for  $P$ , can be organized as a monitor of the form  $\mathcal{M} = (S, s_0, M : S \times \Sigma \rightarrow S)$  for  $P$ . Consequently, the complexity of monitoring  $P$  cannot be smaller than the functional complexity of the partial function  $M : S \times \Sigma \rightarrow S$  corresponding to some “best” monitor  $\mathcal{M}$  for  $P$ ; if there are no additional restrictions, then by “best” monitor we mean the one whose functional complexity of  $M$  is smallest. In particular, if there is no monitor for  $P$  whose transition partial function  $M$  is decidable, then we can safely say that the problem of monitoring  $P$  is undecidable. This discussion leads to the following:

**Pitfall 3.**

*The complexity of monitoring  $P$  is the functional complexity of function  $M$ , where  $\mathcal{M} = (S, s_0, M : S \times \Sigma \rightarrow S)$  is the “best” monitor for  $P$ .*

Since safety properties are precisely the monitorable properties, this appears to be a very natural definition for the complexity of monitoring. While the functional complexity of the monitor function is indeed important because it directly influences the efficiency of monitoring, it is *not* a sufficient measure for the complexity of monitoring. That is because the functional complexity of  $M$  only says how complex  $M$  is in terms of the size of *its input*; it does not say anything about how large the state of the monitor can grow in time. For example, the rewriting-based monitoring algorithm for EREs from [103], whose states are EREs and whose transition is a derivative operation of functional complexity  $O(n^2)$  taking an ERE of size  $n$  into an ERE of size  $O(n^2)$ . It would be very misleading to say that the complexity of monitoring EREs is  $O(n^2)$ , because it may sound much better than it actually is: the  $n^2$  factor accumulates as events are processed. Any monitor for EREs, including the one based on derivatives, eventually requires non-elementary resources (in the size of the ERE) to process a new event.

Therefore, while the complexity of the function  $M$  being executed at each newly received event by a monitor  $\mathcal{M}$  is definitely a necessary and important factor to be considered when defining the complexity of monitoring using  $\mathcal{M}$ , it is not sufficient. One also needs to take into account the size of the input that is being passed to the monitoring function, that is, the size of the monitor state together with the size of the received event. In particular, a monitor storing all the observed trace has unbounded complexity, say  $\infty$ , even though its monitoring function has trivial complexity (e.g., the “event storing” function has linear complexity). More generally, if a property admits no finite-state monitor, then we’d like to say that its monitoring complexity is  $\infty$ : indeed, for any monitor for such a property and for any amount of resources  $R$ , there is some sequence of events that would lead the monitor to a state that needs more than  $R$  resources to be stored or computed. These observations lead us to the following:

**Principle 3.** The complexity of monitoring  $P$  is a function of both the functional complexity of  $M$  and of the size of the states in  $S$ , where  $\mathcal{M} = (S, s_0, M : S \times \Sigma \rightarrow S)$  is an appropriately chosen (“best”) monitor for  $P$ .

We next follow the three principles above and derive our definition for the complexity of monitoring a safety property  $P$ . Before that, let us first define the complexity of monitoring a safety property using a particular monitor for that property, or in other words, let us first define the complexity of a monitor.

During a monitoring session using a monitor, at any moment in time one needs to store at least one state, namely the state that the monitor is currently in. When receiving a new event, the monitor launches its transition function on the current state and the received input. Therefore, the (worst-case) complexity of monitoring with  $\mathcal{M} = (S, s_0, M : S \times \Sigma \rightarrow S)$  could be defined as

$$\max\{FC(M(s, a)) \mid s \in S, a \in \Sigma\},$$

where  $FC(M(s, a))$  is the functional complexity of evaluating  $M$  on state  $s$  and event  $a$ , as a function of the sizes of  $s$  and  $a$ . In other words, the worst-case monitoring complexity of a particular monitor is the maximal functional complexity that its transition function has on any state and any input; this functional complexity is expressed as a function of the size of the pair (state,event). In order for such a definition to make sense formally, one would need to define or axiomatize the size of monitor states and the size of events. Since in order to distinguish  $N$  elements one needs  $\log(N)$  space, we deduce that one needs at least  $\log(|S|)$  space to store the state of the monitor in its worst-case monitoring scenario (each state in  $S$  is reachable).

**Definition 19** *Given a monitor  $\mathcal{M} = (S, s_0, M : S \times \Sigma \rightarrow S)$ , we define the complexity of monitoring  $\mathcal{M}$ , written  $C_{Mon}(\mathcal{M})$ , as the function*

$$FC(M)(\log |S|) : \mathbb{N} \rightarrow \mathbb{N},$$

*which is the “uncurried” version applied on  $\log |S|$  of the worst-case functional complexity  $FC(M) : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$  of the partial function  $M$  as a function of the size of the pair (state,event) being passed to it.*

We assume that the complexity of monitoring a safety property  $P$  is the worst-case complexity of monitoring it using some appropriate, “best” monitor for  $P$ :

$$\min\{\max\{FC(M(s, a)) \mid s \in S, a \in \Sigma\} \mid \mathcal{M} = (S, s_0, M) \in \text{Monitors}(P)\},$$

---

**From safety:** where  $FC(M(s, a))$  is the functional complexity of evaluating  $M$  on state  $s$  and event  $a$ , as a function of the sizes of  $s$  and  $a$ . In other words, we assume that the complexity of monitoring a safety property  $P$  is the worst-case complexity of monitoring it using some appropriate, “best” monitor for  $P$ . The worst-case monitoring complexity of a particular monitor is the maximal functional complexity that its transition function has on any state and any input; this functional complexity is expressed as a function of the size of the pair (state, event). Therefore, in order for such a definition to make sense formally, one would need to define or axiomatize the size of monitor states and the size of events.

---

This gives us the following:

**Definition 20** We let

$$C_{Mon}(P) = \min\{FC(M) \circ \langle \log(|S|), 1_\Sigma \rangle \mid \mathcal{M} = (S, s_0, M) \in \text{Monitors}(P)\}$$

be the complexity of monitoring a safety property  $P$ .

### 4.3 Monitoring Safety Properties is Arbitrarily Hard

We show that the problem of monitoring a safety property can be arbitrarily complex. The previous section tells us that there are as many safety properties as real numbers. Therefore, it is not surprising that some of them can be very hard or impossible to monitor. In this section we formalize this intuitive argument. Our approach is to show that we can associate a safety property  $P_S$  to any set of natural numbers  $S$ , such that monitoring that safety property is as hard as checking membership of arbitrary natural numbers to  $S$ . The result then follows from the fact that checking memberships of natural numbers to sets of natural numbers is a problem that can be arbitrarily complex.

Theorem 2 indirectly says that we can associate a persistent safety property to any set of natural numbers (sets of natural numbers are in a bijective correspondence with the real numbers). However, it is not clear how that safety property looks and neither how to monitor it. We next give a more concrete mapping from sets of natural numbers to (persistent) safety properties and show that monitoring the property is equivalent to

testing membership to the set. It suffices to assume that  $\Sigma$  contains only two elements, say  $\Sigma = \{0, 1\}$ .

**Definition 21** Let  $P_- : \mathcal{P}(\mathbb{N}) \rightarrow \text{PersistentSafety}^*$  be the mapping defined as follows: for any  $S \subseteq \mathbb{N}$ , let  $P_S$  be the set  $1^* \cup \{1^k 0 \mid k \in S\} \cdot \{0, 1\}^*$ .

It is easy to see that  $P_S$  is a persistent safety property over finite traces. Also, it is easy to see that the bijection in the proof of Theorem 3 associates to  $P_S$  the safety property over infinite traces  $1^\omega \cup \{1^k 0 \mid k \in S\} \cdot \{0, 1\}^\omega$ .

Let us now investigate the problem of monitoring  $P_S$ .

**Proposition 11** For any  $S \subseteq \mathbb{N}$ , monitoring  $P_S$  is equivalent to deciding membership of natural numbers to  $S$ .

**Proof:** If  $M_S$  is an oracle deciding membership of natural numbers to  $S$ , that is, if  $M_S(n)$  is true iff  $n \in S$ , then one can build a monitor for  $P_S$  as follows: for a given trace, incrementally read and count the number of prefix 1's; if no 0 is ever observed then monitor indefinitely without reporting any violation; when a first 0 is observed, if any, ask if  $M(k)$ , where  $k$  is the number of 1's observed; if  $M(k)$  is false, then report violation; if  $M(k)$  is true, then continue monitoring indefinitely and never report violation. It is clear that this is indeed a monitor for  $P_S$ .

Conversely, if we had any monitor for  $P_S$  then we could build a decision procedure for membership to  $S$  as follows: given  $k \in \mathbb{N}$ , send to the monitor a sequence of  $k$  ones followed by a 0; if the monitor reports violation then deduce that  $k \notin S$ ; if the monitor does not report violation, then deduce that  $k \in S$ . It is clear that this is a decision procedure for membership to  $S$ .

The proof works for both persistent safety properties over finite traces and for safety properties over infinite traces.  $\square$

The claim in the title of this section follows now from the fact that the set  $S$  of natural numbers can be chosen so that its membership problem is arbitrarily complex. For example, since there are as many subsets of natural numbers as real numbers while there are only as many Turing machines as natural numbers, it follows that there are many (exponentially) more sets of natural numbers that are not recognized by Turing machines than those that are. In particular, there are sets of natural numbers corresponding to any degree in the arithmetic hierarchy, i.e., to predicates  $A(k)$  of the form  $(Q_1 k_1)(Q_2 k_2) \cdots (Q_n k_n) R(k, k_1, k_2, \dots, k_n)$ , where  $Q_1, Q_2, \dots, Q_n$  are alternating (universal or existential) quantifiers and  $R$  is a recursive/decidable

relation: for  $A$  such a predicate, let  $S_A$  be the set of natural numbers  $\{k \mid A(k)\}$ . Recall that if  $Q_1$  is  $\forall$  then  $A$  is called a  $\Pi_n$  property, while if  $Q_1$  is  $\exists$  then  $A$  is called a  $\Sigma_n$  property. In particular,  $\Sigma_0 = \Pi_0$  and they contain precisely the recursive/decidable properties,  $\Sigma_1$  contains precisely the class of recursively enumerable problems,  $\Pi_1$  contains precisely the co-recursively enumerable problems, etc.; a standard  $\Pi_2$  problem is TOTALITY: given  $k \in \mathbb{N}$ , is it true that Turing machine with Gödel number  $k$  terminates on all inputs? Since each level in the arithmetic hierarchy contains problems strictly harder than problems on the previous layer (because  $\Sigma_n \cup \Pi_n \subsetneq \Sigma_{n+1} \cap \Pi_{n+1}$ ), the arithmetic hierarchy gives us a universe of safety properties whose monitoring can be arbitrarily hard.

Within the decidable fragment, as expected, monitoring safety properties can also have any complexity. Indeed, pick for example any NP-complete problem and let  $S$  be the set of inputs (coded as natural numbers) for which the problem has a positive answer; then, as explained in the proof of Proposition 11, monitoring  $P_S$  against input  $1^k0$  is equivalent to deciding membership of  $k$  to  $S$ , which is further equivalent to answering the NP-complete problem on input  $k$ . Of course, in practice a particular (implementation of a) monitor can be more complex than the corresponding membership problem; for example, monitors corresponding to NP-complete problems are most likely exponential. Also, note that a monitor for  $P_S$  needs not necessarily do its complex computation on an input  $1^k0$  when it encounters the 0. It can perform intermediate computations as it reads the prefix 1's and thus pay a lesser computational price when the 0 is encountered. What Proposition 11 says is that the *total* complexity to process the input  $1^k0$  can be no lower than the complexity of checking whether  $k \in S$ .

## 4.4 Canonical Monitors

We conclude this section with an alternative definition of a monitor, called *canonical monitor*, which is more compact than our previous definition and which appears to be sufficient to capture any safety property. We do not make any use of this alternative definition in this paper, but it may serve as a basis for further foundational endeavors in this area.

The set of states  $S$  of a monitor  $(S, s_0, M : S \times \Sigma \rightarrow S)$  are typically enumerable, so they can be very well replaced with natural numbers. Moreover, the initial state  $s_0$  can be encoded, by convention, as the first natural number, 0. A monitor then becomes nothing but a partial function  $\mathbb{N} \times \Sigma \rightarrow \mathbb{N}$ . We

therefore rightfully call these particular monitors *canonical*:

**Definition 22** A canonical  $\Sigma$ -monitor is a partial function  $\mathcal{N} : \mathbb{N} \times \Sigma \rightarrow \mathbb{N}$ . Let  $\mathcal{S}_{\mathcal{N}} = \{n \mid (\exists w) \mathcal{N}(0, w) = n\}$  be the states of  $\mathcal{N}$ . As before, let

- $\mathcal{L}^*(\mathcal{N}) = \{w \in \Sigma^* \mid \mathcal{N}(0, w) \downarrow\}$ ,
- $\mathcal{L}^\omega(\mathcal{N}) = \{u \in \Sigma^\omega \mid \mathcal{N}(0, w) \downarrow \text{ for all } w \in \text{prefixes}(u)\}$ , and
- $\mathcal{L}^{*,\omega}(\mathcal{N}) = \mathcal{L}^*(\mathcal{N}) \cup \mathcal{L}^\omega(\mathcal{N})$ .

Although the set of states  $S$  in a monitor  $(S, s_0, M : S \times \Sigma \rightarrow S)$  is allowed to have any cardinal while the states in canonical monitors are restricted to natural numbers, it turns out that canonical monitors can in fact express all monitorable properties:

**Proposition 12** A property  $P \subseteq \Sigma^*$  (resp.  $P \subseteq \Sigma^\omega$ , resp.  $P \subseteq \Sigma^* \cup \Sigma^\omega$ ) is monitorable iff there is some canonical monitor  $\mathcal{N}$  such that  $P = \mathcal{L}^*(\mathcal{N})$  (resp.  $P = \mathcal{L}^\omega(\mathcal{N})$ , resp.  $P = \mathcal{L}^{*,\omega}(\mathcal{N})$ ).

**Proof:** Since any canonical monitor is a monitor, it follows that any property specifiable by a canonical monitor is indeed monitorable. For the converse, let  $P$  be a property monitorable by some monitor  $\mathcal{M} = (S, s_0, M : S \times \Sigma \rightarrow S)$ . Since  $|\Sigma| \leq \aleph_0$ , we can enumerate all the states of  $\mathcal{M}$  that can be reached from  $s_0$  with its transition function  $M$ . There are many different ways to do this (e.g., in breadth-first order, in depth-first order, etc.), but these are all ultimately irrelevant. If we let  $S^r = \{s_0, s_1, s_2, \dots\}$  denote the resulting set of reachable states, then it is easy to first note that the monitor  $\mathcal{M}^r = (S^r, s_0, M : S^r \times \Sigma \rightarrow S^r)$  specifies the same property  $P$  as  $\mathcal{M}$ , and second note that  $\mathcal{M}^r$  specifies the same property as the canonical monitor  $\mathcal{N} : \mathbb{N} \times \Sigma \rightarrow \mathbb{N}$  defined by  $\mathcal{N}(i, a) = j$  iff  $M(s_i, a) = s_j$ .  $\square$

## Exercises

**Exercise 7** Define a canonical monitor for the property

“A file can only be accessed if it is open.”

That is, the file can only be accessed if it was opened at some moment in the past and it was not closed since then. Suppose  $\Sigma$  consists of the events/actions  $\{o, a, c\}$ , where  $o$  stands for “file open”,  $a$  for “file access”, and  $c$  for “file close”.



## Chapter 5

# Event/Trace Observation



## Chapter 6

# Monitor Synthesis: Finite State Machines (FSM)

---

*discuss also DFA, NFA, using NFA as monitor, REs, RE2NFA, derivatives*

---

### 6.1 Binary Transition Trees (BTT)

#### 6.1.1 Multi-Transition and Binary Transition Tree Finite State Machines

To keep the runtime overhead of monitors low, it is crucial to do as little computation as possible in order to proceed to the next state. In the sequel we assume that finite state monitors are desired to efficiently change their state (when a new event is received) to one of possible states  $s_1, s_2, \dots, s_n$ , under the knowledge that a transition to each such state is enabled deterministically by some Boolean formula,  $p_1, p_2, \dots, p_n$ , respectively, on atomic state predicates.

**Definition 23** *Let  $S$  be a set whose elements are called states, and let  $A$  be another set, whose elements are called atomic predicates. Let  $\{s_1, s_2, \dots, s_n\} \subseteq S$  and let  $p_1, p_2, \dots, p_n$  be propositions over atoms in  $A$  (using the usual Boolean combinators presented in Subsection 2.1.2), with the property that exactly one of them is true at any moment, that is,  $p_1 \vee p_2 \vee \dots \vee p_n$  holds and for any distinct  $p_i, p_j$ , it is the case that  $p_i \rightarrow \neg p_j$  holds. Then we call the  $n$ -tuple  $[p_1?s_1, p_2?s_2, \dots, p_n?s_n]$  a multi-transition (MT) over states*

$S$  and atomic predicates (or simply atoms)  $A$ . Let  $MT(S, E)$  be the set of multi-transitions over states  $S$  and atoms  $A$ .

Since the rest of the paper is concerned with rather theoretical results, from now on we use mathematical symbols for the Boolean operators instead of ASCII symbols like in Subsection 2.1.2. The intuition underlying multi-transitions is straightforward: depending on which of the propositions  $p_1, p_2, \dots, p_n$  is true at a given moment, a corresponding transition to exactly one of the states  $s_1, s_2, \dots, s_n$  will take place. Formally,

**Definition 24** Maps  $\theta : A \rightarrow \{\text{true}, \text{false}\}$  are called events from now on. With the notation above, given an event  $\theta$ , we define a map  $\theta_{MT} : MT(S, A) \rightarrow S$  as  $\theta_{MT}([p_1?s_1, p_2?s_2, \dots, p_n?s_n]) = s_i$ , where  $\theta(p_i) = \text{true}$ ; notice that  $\theta(p_j) = \text{false}$  for any  $1 \leq j \neq i \leq n$ .

**Definition 25** Given an event  $\theta : A \rightarrow \{\text{true}, \text{false}\}$ , let  $e_\theta$  denote the list of atomic predicates  $a$  with  $\theta(a) = \text{true}$ . There is an obvious correspondence between events as maps  $A \rightarrow \{\text{true}, \text{false}\}$  and events as lists of atomic predicates, which justifies our implementation of events in Subsection 2.1.2. We take the liberty to use either the map or the list notation for events from now on in the paper. We let  $\mathcal{E}$  denote the set of all events and call lists, or words,  $e_{\theta_1} \dots e_{\theta_n} \in \mathcal{E}^*$  (finite) traces; this is also consistent with our mechanical Maude ASCII notation in Subsection 2.1.2, except that we prefer not to separate events by commas in traces from now on, to avoid mathematical notational conflicts.

We next define binary transition trees.

**Definition 26** Under the same notations as in the previous definition, a binary transition tree (BTT) over states  $S$  and atoms  $A$  is a term over syntax

$$BTT ::= S \mid (A ? BTT : BTT).$$

We let  $BTT(S, A)$  denote the set of binary transition trees over states  $S$  and atoms  $A$ . Given an event  $\theta : A \rightarrow \{\text{true}, \text{false}\}$ , we define a map  $\theta_{BTT} : BTT(S, A) \rightarrow S$  inductively as follows:

$$\begin{aligned} \theta_{BTT}(s) &= s \text{ for any } s \in S, \\ \theta_{BTT}(a ? b_1 : b_2) &= \theta_{BTT}(b_1) \text{ if } \theta(a) = \text{true}, \text{ and} \\ \theta_{BTT}(a ? b_1 : b_2) &= \theta_{BTT}(b_2) \text{ if } \theta(a) = \text{false}. \end{aligned}$$

A binary transition tree  $b$  in  $BTT(S, A)$  is said to implement a multi-transition  $t$  in  $MT(S, A)$  if and only if  $\theta_{BTT}(b) = \theta_{MT}(t)$  for any map  $\theta : A \rightarrow \{true, false\}$ .

BTTs generalize BDDs [14], which can be obtained by taking  $S = \{true, false\}$ . As an example of BTT,  $a1 ? a2 ? s1 : a3 ? false : s2 : a3 ? s2 : true$  says “eval  $a_1$ ; if  $a_1$  then (eval  $a_2$ ; if  $a_2$  then go to state  $s1$  else (eval  $a_3$ ; if  $a_3$  then go to state  $false$  else go to state  $s2$ )) else (eval  $a_3$ ; if  $a_3$  then go to state  $s2$  else go to state  $true$ )”. Note that  $true$  and  $false$  are just some special states. Depending on the application they can have different meanings, but in our applications  $true$  typically means that the monitoring requirement has been fulfilled, while  $false$  means that it has been violated. It is often convenient to represent BTTs graphically, such as the one in Figure 6.1.

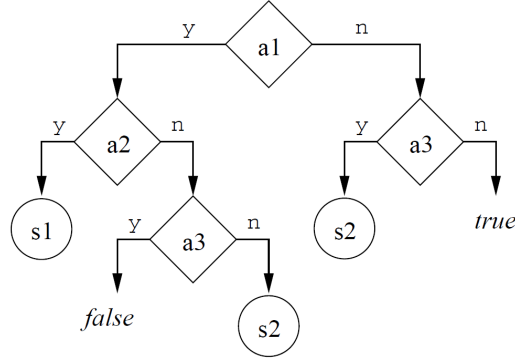


Figure 6.1: Graphical representation for the BTT  $a1 ? a2 ? s1 : a3 ? false : s2 : a3 ? s2 : true$ .

**Definition 27** A multi-transition finite state machine (*MT-FSM*) is a triple  $(S, A, \mu)$ , where  $S$  is a set of states,  $A$  is a set of atomic predicates, and  $\mu$  is a map from  $S - \{true, false\}$  to  $MT(S, A)$ . If  $S$  contains  $true$  and/or  $false$ , then  $\mu(true) = [true?true]$  and/or  $\mu(false) = [true?false]$ , respectively. A binary transition tree finite state machine (*BTT-FSM*) is a triple  $(S, A, \beta)$ , where  $S$  and  $A$  are like before and  $\beta$  is a map from  $S - \{true, false\}$  to  $BTT(S, A)$ . If  $S$  contains  $true$  and/or  $false$ , then it is the case that  $\beta(true) = [true]$  and/or  $\beta(false) = [false]$ , respectively. For an event  $\theta : A \rightarrow \{true, false\}$  in any of these machines, we let  $s \xrightarrow{\theta} s'$  denote the fact that  $\theta_{MT}(\mu(s)) = s'$  or

$\theta_{BTT}(\beta(s)) = s'$ , respectively. We take the liberty to call  $s \xrightarrow{\theta} s'$  a transition. Also, we may write  $s_1 \xrightarrow{\theta_1} s_2 \xrightarrow{\theta_2} \dots s_n \xrightarrow{\theta_n} s_{n+1}$  and call it a sequence of transitions whenever  $s_1 \xrightarrow{\theta_1} s_2, \dots, s_n \xrightarrow{\theta_n} s_{n+1}$  are transitions.

Note that  $S$  is not required to contain the special states *true* and *false*, but if it contains them, these states transit only to themselves. The finite state machine notions above are intended to abstract the intuitive concept of a monitor. The states in  $S$  are monitor states, and some of them can trigger side effects in practice, such as messages sent or reported to users, actions taken, feedback sent to the monitored program for guidance purposes, etc.

**Definition 28** *If  $true \in S$  then a sequence of transitions  $s_1 \xrightarrow{\theta_1} s_2 \xrightarrow{\theta_2} \dots s_n \xrightarrow{\theta_n} true$  is called a validating sequence (for  $s_1$ ); if  $false \in S$  then  $s_1 \xrightarrow{\theta_1} s_2 \xrightarrow{\theta_2} \dots s_n \xrightarrow{\theta_n} false$  is called an invalidating sequence (for  $s_1$ ). These notions naturally extend to corresponding finite traces of events  $e_{\theta_1} \dots e_{\theta_n}$ .*

BTT-FSMs can and should be thought of as efficient “implementations” of MT-FSMs, in the sense that they implement multi-transitions as binary transition trees. These allow one to reduce the amount of computation in a monitor implementing a BTT-FSM to only evaluate at most all the atomic predicates in a given state in order to decide to which state to go next; however, it will only very rarely be the case that *all* the atomic predicates need to be evaluated.

A natural question is why one should bother at all then with defining and investigating MT-FSMs. Indeed, what one really looks for in the context of FSM monitoring is efficient BTT-FSMs. However, MT-FSMs are nevertheless an important intermediate concept as it will become clearer later in the paper. This is because they have the nice property of *state mergeability*, which allows one to elegantly generate MT-FSMs from logical formulae. By state mergeability we mean the following. Suppose that during a monitor generation process, such as that for LTL that will be presented in Subsection 8.4.1, one proves that states  $s$  and  $s'$  are “logically equivalent” (for now, equivalence can be interpreted intuitively: they have the same behavior), and that  $s$  and  $s'$  have the multi-transitions  $[p_1?s_1, p_2?s_2, \dots, p_n?s_n]$  and  $[p'_1?s'_1, p'_2?s'_2, \dots, p'_n?s'_n]$ . Then we can merge  $s$  and  $s'$  into one state whose multi-transition is  $Merge([p_1?s_1, p_2?s_2, \dots, p_n?s_n], [p'_1?s'_1, p'_2?s'_2, \dots, p'_n?s'_n])$ , defined as follows:

$Merge([p_1?s_1, p_2?s_2, \dots, p_n?s_n], [p'_1?s'_1, p'_2?s'_2, \dots, p'_{n'}?s'_{n'}])$   
contains all choices  $p?s''$ , where  $s''$  is a state in  $\{s_1, s_2, \dots, s_n\} \cup \{s'_1, s'_2, \dots, s'_{n'}\}$  and

- $p$  is  $p_i$  when  $s'' = s_i$  for some  $1 \leq i \leq n$  and  $s'' \neq s'_{i'}$  for all  $1 \leq i' \leq n'$ , or
- $p$  is  $p'_{i'}$  when  $s'' = s'_{i'}$  for some  $1 \leq i' \leq n'$  and  $s'' \neq s_i$  for all  $1 \leq i \leq n$ , or
- $p$  is  $p_i \vee p'_{i'}$  when  $s'' = s_i$  for some  $1 \leq i \leq n$  and  $s'' = s'_{i'}$  for some  $1 \leq i' \leq n'$ .

It is easy to see that this multi-transition merging operation is well defined, that is,

**Proposition 13**  *$Merge([p_1?s_1, p_2?s_2, \dots, p_n?s_n], [p'_1?s'_1, p'_2?s'_2, \dots, p'_{n'}?s'_{n'}])$  is a well-formed multi-transition whenever both  $[p_1?s_1, p_2?s_2, \dots, p_n?s_n]$  and  $[p'_1?s'_1, p'_2?s'_2, \dots, p'_{n'}?s'_{n'}]$  are well-formed multi-transitions. Therefore, MERGE can be seen as a function  $\mathcal{P}_f(MT(S, A)) \rightarrow MT(S, A)$ , where  $\mathcal{P}_f$  is the finite powerset operator.*

There are situations when a definite answer, *true* or *false* is desired at the end of the monitoring session, as it will be in the case of LTL. As explained in Section 8.1, the intuition for the last event in an execution trace is that the trace is infinite and stationary in that last event. This seems to be the best and simplest assumption about future when a monitoring session is ended.

---

*Discussion about variants of finite-trace LTL needed.*

---

For such situations, we enrich our definitions of MT-FSM and BTT-FSM with support for terminal events:

**Definition 29** *A terminating multi-transition finite state machine (abbreviated MT-FSM\*) is a tuple  $(S, A, \mu, \mu^*)$ , where  $(S, A, \mu)$  is an MT-FSM and  $\mu^*$  is a map from  $S - \{true, false\}$  to  $MT(\{true, false\}, A)$ . A terminating binary transition tree finite state machine (BTT-FSM\*) is a tuple  $(S, A, \beta, \beta^*)$ , where  $(S, A, \beta)$  is a BTT-FSM and  $\beta^*$  is a map from  $S - \{true, false\}$  to  $BTT(\{true, false\}, A)$ . For a given event  $\theta : A \rightarrow \{true, false\}$  in any of these finite state machines, we let  $s \xrightarrow{\theta^*} true$  (or *false*) denote the fact*

that  $\theta_{MT}(\mu^*(s)) = \text{true}$  (or  $\text{false}$ ) or  $\theta_{BTT}(\beta^*(s)) = \text{true}$  (or  $\text{false}$ ), respectively. We call  $s \xrightarrow{\theta^*} \text{true}$  (or  $\text{false}$ ) a terminating transition. A sequence  $s_1 \xrightarrow{\theta_1} s_2 \xrightarrow{\theta_2} \dots s_n \xrightarrow{\theta_n^*} \text{true}$  is called an accepting sequence (for  $s_1$ ) and a sequence  $s_1 \xrightarrow{\theta_1} s_2 \xrightarrow{\theta_2} \dots s_n \xrightarrow{\theta_n^*} \text{false}$  is called a rejecting sequence (for  $s_1$ ). These notions also naturally extend to corresponding finite traces of events  $e_{\theta_1} \dots e_{\theta_n}$ .

Languages can be associated to states in MT-FSM\*s or BTT-FSM\*s as finite words of events.

**Definition 30** Given a state  $s \in S$  in an MT-FSM\* or in a BTT-FSM\*  $M$ , we let  $\mathcal{L}_M(s)$  denote the set of finite traces  $e_{\theta_1} \dots e_{\theta_n}$  with the property that  $s \xrightarrow{\theta_1} \dots \xrightarrow{\theta_n^*} \text{true}$  is an accepting sequence in  $M$ . If a state  $s_0 \in S$  is desired to be initial, then we write it at the end of the tuple, such as  $(S, A, \mu, \mu^*, s_0)$  or  $(S, A, \beta, \beta^*, s_0)$ , and let  $\mathcal{L}_M$  denote  $\mathcal{L}_M(s_0)$ . If  $s_0 \xrightarrow{\theta_1} s_1 \dots \xrightarrow{\theta_n} s_n$  is a sequence of transitions from the initial state, then  $e_{\theta_1} \dots e_{\theta_n}$  is called a valid prefix if and only if  $e_{\theta_1} \dots e_{\theta_n} t \in \mathcal{L}_M$  for any (empty or not) trace  $t$ , and it is called an invalid prefix if and only if  $e_{\theta_1} \dots e_{\theta_n} t \notin \mathcal{L}_M$  for any trace  $t$ .

The following is immediate.

**Proposition 14** If  $\text{true} \in S$  then  $\mathcal{L}_M(\text{true}) = \mathcal{E}^*$ , and if  $\text{false} \in S$  then  $\mathcal{L}_M(\text{false}) = \emptyset$ . If  $s \xrightarrow{\theta} s'$  in  $M$  then  $\mathcal{L}_M(s') = \{t \mid e_{\theta} t \in \mathcal{L}_M(s)\}$ ; more generally, if  $s \xrightarrow{\theta_1} s_1 \dots \xrightarrow{\theta_n} s_n$  is a sequence of transitions in  $M$  then  $\mathcal{L}_M(s_n) = \{t \mid e_{\theta_1} \dots e_{\theta_n} t \in \mathcal{L}_M(s)\}$ . In particular, if  $s = s_0$  then  $e_{\theta_1} \dots e_{\theta_n}$  is a valid prefix if and only if  $\mathcal{L}_M(s_n) = \mathcal{E}^*$ , and it is an invalid prefix if and only if  $\mathcal{L}_M(s_n) = \emptyset$ .

### 6.1.2 From MT-FSMs to BTT-FSMs

Supposing that one has encoded a logical requirement into an MT-FSM (we shall see how to do it for LTL in the next subsection), the next important step is to generate an efficient equivalent BTT-FSM. In the worst possible case one just has to evaluate all the atomic predicates in order to proceed to the next state of a BTT-FSM, so they are preferred to MT-FSMs. What one needs to do is to develop a procedure that takes a multi-transition and returns a BTT. More BTTs can encode the same multi-transition, so one needs to develop some criteria to select the better ones. A natural selection

criterion would be to minimize the average amount of computation. For example, if all atomic predicates are equally probable to hold and if an atomic predicate is very expensive to evaluate, then one would select that BTT that places the expensive predicate as deeply as possible, so its evaluation is delayed as much as possible. Based on the above, we believe that the following is an important theoretical problem in runtime monitoring:

**Problem:** Optimal BTT

**Input:** A set of atomic predicates  $a_1, \dots, a_k$  that hold with probabilities  $\pi_1, \dots, \pi_k$  and have costs  $c_1, \dots, c_k$ , respectively, and a multi-transition  $p_1?s_1, \dots, p_n?s_n$  where  $p_1, \dots, p_n$  are Boolean formulae over  $a_1, \dots, a_k$ .

**Output:** A BTT implementing the multi-transition that probabilistically minimizes the amount of computation to decide the next state.

The probabilities and the costs in the problem above can be estimated either by static analysis of the source code of the program, or dynamically by first running and measuring the program several times, or by combinations of those. We do not know how to solve this interesting problem yet, but we conjecture the following result that we currently use in our implementation:

**Conjecture 7** *If  $\pi_1 = \dots = \pi_k$  and  $c_1 = \dots = c_k$  then the solution to the problem above is the BTT of minimal size.*

Our current multi-transition to BTT algorithm is exponential in the number of atomic predicates; it simply enumerates all possible BTTs recursively and then selects the one of minimal size. Generating the BTTs takes significantly less time than generating the MT-FSM, so we do not regard it as a practical problem yet. However, it seems to be a challenging theoretical problem.

**Example 6** *Consider again the traffic light controller safety property stating that “after green comes yellow”, which was written as  $\square(\text{green} \rightarrow (!\text{red} \text{ U } \text{yellow}))$  using LTL notation. Since more than one light can be lit at any moment, one should be very careful when expressing this safety property as an MT-FSM or a BTT-FSM.*

*Let us first express it as an MT-FSM. We need two states, one for the case in which green has not triggered yet the  $!\text{red} \text{ U } \text{yellow}$  part and*

State	MT for non-terminal events	MT for terminal events
1	<pre>[     yellow \/ !green ? 1,     !yellow /\ green /\ !red ? 2,     !yellow /\ green /\ red ? false ]</pre>	<pre>[ yellow \/ !green ? true,   !yellow /\ green ? false ]</pre>
2	<pre>[     yellow ? 1,     !yellow /\ !red ? 2,     !yellow /\ red ? false ]</pre>	<pre>[ yellow ? true,   !yellow ? false ]</pre>

Figure 6.2: MT-FSM for the formula  $[(\text{green} \rightarrow \text{!red} \cup \text{yellow})]$ .

another for the case when it has. The condition to stay in state 1 is then `yellow \/ !green` and the condition to move to state 2 is `!yellow /\ green /\ !red`. If both a `green` and a `red` are seen then the machine should move to state false. The condition to move from state 2 back to state 1 is `yellow`, while the condition to stay in state 2 is `!yellow /\ !red`; `!yellow /\ red` should also move the machine in its false state. If the event is terminal then a `yellow` would make the reported answer true, i.e., the observed trace is an accepting sequence; if `yellow` is not true, then in state 1 the answer should be the opposite of `green`, while in state 2 the answer should be simply false. This MT-FSM is shown in Figure 6.2.

The interesting aspect of our FSMs is that not all the atomic predicates need to always be evaluated. For example, in state 2 of this MT-FSM the predicate `green` is not needed. A BTT-FSM further reduces the amount of predicates to be evaluated, by enforcing an “evaluate-by-need” policy. Figure 6.3 shows a BTT-FSM implementing the MT-FSM in Figure 6.2.

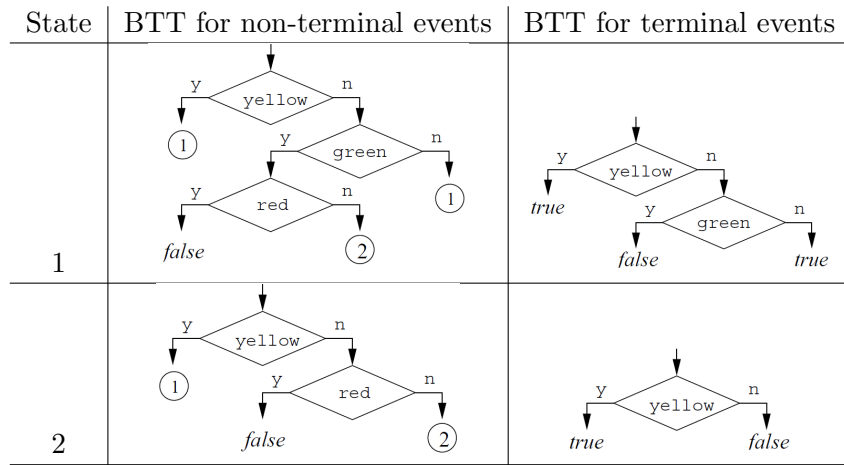


Figure 6.3: A BTT-FSM for the formula  $\Box(\text{green} \rightarrow !\text{red} \cup \text{yellow})$ .



## Chapter 7

# Monitor Synthesis: Extended Regular Expressions (ERE)

---

*bad vs. good prefixes — cite Vardi*

---

### 7.1 Monitoring ERE Safety Needs Non-Elementary Space

Extended regular expressions (EREs) add complementation ( $\neg$ ) to regular expressions (REs). Complementation can be handy when defining safety properties, because it allows us to say both what should never happen as well as what could happen. In particular, in the context of monitoring EREs, one can switch between the expression of bad prefixes of a safety property and that of good prefixes by just applying a complement operator.

In this section we show that any monitor for safety properties expressed as EREs requires non-elementary space. More precisely, for a given  $n \in \mathbb{N}$ , we build an ERE of size  $O(n^3)$  such that its language is prefix closed and any monitor for its language requires space

$$2^{2^{\cdot^{2^n}}} \quad \Bigg] \quad n \text{ nested power of 2 operations}$$

---

*discuss that this also implies non-elementary time complexity*

---

### 7.1.1 Discussion and Relevance of the Lower-Bound Result

Since regular expressions (REs) and deterministic (DFA) and non-deterministic (NFA) finite-state automata are enumerable structures while the set of safety properties is in bijection with the set of real numbers, there are many safety properties that cannot be expressed using REs or automata (or any other formalism whose objects are enumerable). Nevertheless, there are many safety properties of interest that can be expressed as REs or automata. Safety properties over finite or infinite traces can be expressed as REs in at least two different ways:

1. Use an RE to express the language of its bad prefixes; or
2. Use an RE to express the language of its (good) prefixes.

In the first case, the RE captures the finite-trace behaviors that should never happen, while in the second the RE captures the ones that could possibly happen.

Obviously, not all REs correspond to safety properties in one or the other of the two cases above. For example, the RE  $(0 \cdot 1)^+$  cannot express the bad prefixes of a safety property, because “01” is a bad prefix while “010” is not. In order for an RE to express the bad prefixes of a safety property, it should have the property that once  $w$  is in its language, then all  $ww'$  for any  $w'$  should also be in its language. The RE  $(0 \cdot 1)^+$  cannot express the good prefixes of a safety property either: “0101” is a good prefix while 010 is not. The language of an RE must be prefix closed in order to express the good prefixes of a safety property.

In both cases above, monitoring the safety property can be done very efficiently (linearly in the size of the RE, both space-wise and time-wise) by first translating its corresponding RE into an NFA, for example using a technique such as Thompson’s [127], and then doing one of the following:

1. In the first case, simulate the NFA-to-DFA construction on the fly as events are received. The state of the monitor is therefore a set of states of the NFA. At start, that set contains only the initial state of the NFA. For each new event, construct the next set by collecting all the NFA states that can be reached via the received event from any of the existing states in the set. If a final state is reached then report violation of the property: bad prefix found. Since the final states in the NFA symbolize the reach of a bad prefix and since bad prefixes have “no future” in a safety property, the NFA associated to the original RA

can be optimized (in case it is not already optimal directly from its construction) by removing any edges out of its final states.

2. In the second case, the monitor works the same way as in the first case, but checking at each time that the monitor state (also a set of NFA states) contains at least one final state of the NFA; if that is not the case, then report violation: prefix found which is not good. If one is willing to pay the exponential price and determinize the NFA of good prefixes, then one can further optimize the resulting DFA (in case it is not already optimized by construction) by collapsing all its non-final states into a “dead-end” state: indeed, the reachability of a non-final state signifies the reachability of a bad prefix, which has “no future”. It is not clear whether or how to optimize the NFA of good prefixes using the additional info that it is a safety property.

Extended regular expressions (EREs) add complementation ( $\neg$ ) to REs. Meyer and Stockmeyer [124, 125] showed that EREs can express languages non-elementarily more compactly than REs. In other words, for any constant  $k \geq 1$ , one can find EREs of large enough size  $n \in \mathbb{N}$  for which there is no RE having the same language of size less than  $2^{2^{\cdot^{2^n}}}$ , with  $k$  nested power operations. Meyer and Stockmeyer [124, 125] showed that several other problems concerning EREs are also non-elementary, including: the equivalence of EREs, the emptiness of the language of an ERE (and implicitly the emptiness of the complement of the language of an ERE), the automata generation (NFA or DFA) from an ERE, etc. Note that it is straightforward to generate potentially non-elementarily large automata from EREs. All one needs to do is to iteratively apply NFA-to-DFA transformations for EREs under complement operators and then complement the resulting DFAs (by complementing their final states). Since each NFA-to-DFA transformation may lead to an exponential explosion on the number of states and since the ERE can have arbitrarily many nested complement operators, the resulting NFA or DFA can be non-elementarily larger than the original ERE.

As already mentioned above, if we allow complementation then we can easily switch from an expression defining the bad prefixes of a safety property to one defining its good prefixes, and backwards, by applying a complement operator. Therefore, from here on, when we say that an ERE expresses a safety property, without any loss of generality we assume that it defines the good prefixes of the safety property; in particular, we assume that its language is prefix closed. Clearly, if one can afford to generate an automaton

from the ERE expressing a safety property, then one can and probably should use that automaton as a monitor for the safety property. However, since such an automaton can be enormous compared to the size of the original ERE, a natural question to ask is whether one can generate monitors for safety properties expressed as EREs that need less than non-elementary space in the size of the original ERE. We next give a negative answer to this question.

Let us first discuss our subsequent lower bound result from a more conceptual perspective. Notice that *synchronous monitoring*, that is, the monitoring process where an error is detected as soon as it appears, is harder than checking for satisfiability (or non-emptiness); indeed, if a formula in a particular formalism is not satisfiable (or it has an empty language), then a synchronous monitor should detect that before the first event is observed. We refer the interested reader to [105] for a discussion on various types of monitoring, including synchronous versus asynchronous monitoring. Synchronous monitors need to either directly (e.g., by accumulating logical constraints while verifying their consistency) or indirectly (e.g., by generating statically an automaton or a structure containing all possible future behaviors) check for satisfiability (or non-emptiness) of the remaining requirements as events are observed online. Unfortunately, this is an expensive process that may be desired to be avoided, at the expense of delaying the detection of violations. For example, the rewriting based monitoring approach for LTL in [105] delays the detection of violations of LTL formulae of the form “(next  $\varphi$ )” and “(next  $\neg\varphi$ )” for one more event, to avoid invoking an expensive satisfiability checker for LTL but to instead invoke a propositional satisfiability checker which is less expensive in practice; this is closely related to the notion of “informative prefixes” in [83] that “tell all the story”. Our subsequent lower-bound result states that any monitor for safety properties expressed using EREs, *synchronous or asynchronous*, requires non-elementary space.

Let us now clarify that our lower bound result is not a consequence of the lower-bound result by Meyer and Stockmeyer in [124] (see [125] for a proof of that result). A first reason is that neither the ERE constructed in [125, 124] nor its complement is prefix closed. Indeed, we here focus on a subset of EREs, rather than arbitrary EREs, namely those whose languages are prefix closed, so they express good prefixes of safety properties. Supposing that one could modify the “hard” ERE in [125, 124] whose complement non-emptiness requires non-elementary space into one whose language is prefix-closed and whose size is linear in the size of the original one, the fact that synchronous

monitoring of EREs is harder than checking for emptiness does not necessarily imply that monitoring that hard ERE requires non-elementary space. In fact, monitoring that particular ERE requires constant time to process each event, because it is equivalent with an automaton of one state – it takes, however, non-elementary space to compute that one state automaton. What it says is therefore that the *initialization step of ERE-safety synchronous monitoring* requires, in the worst case, non-elementary space.

Interestingly, if one could modify the results in [125, 124] to hold for prefix-closed EREs, including especially the result stating that finding an RE equivalent to an ERE is a non-elementary problem, then one could show that *synchronous monitoring* of safety properties expressed as EREs requires non-elementary space! Indeed, supposing that one had for any ERE a monitor that takes only elementary space in its worst-case monitoring scenario, then one could use that monitor to generate a DFA for the ERE as follows: start with the initial state of the monitor and then discover and store new states of the monitor by feeding arbitrary (but finite in number) events to each state of the monitor until no new state is discovered. This closure operation takes as much time and space as the number of states the monitor reach; since by assumption the monitor needs “only” elementary space to store its state in any scenario, we deduce that the obtained automaton has size elementary in the size of the ERE (and it also takes elementary time and space to generate it). In other words, we could find an elementary algorithm to associate to any (prefix closed) ERE an equivalent RE, contradicting the non-elementary lower bound in [125, 124] (again, supposing that the lower-bound results in [125, 124] could be modified for prefix-closed EREs).

Unfortunately, it is not that clear how to reduce the non-elementary problems in [125, 124] to *asynchronous monitoring*, and thus to conclude that asynchronous monitoring also requires non-elementary space. That is because a “smart” asynchronous monitor may in principle collapse states (when regarded as an automaton as in the construction above) in rather unexpected ways, just because it “knows” that eventually an error may be reported anyway if observation continues indefinitely from that state on; in other words, states with the property that “eventually violation detected in the future” may be collapsed as equivalent by an asynchronous monitor. This way, the DFA extracted from a monitor for a safety property expressed as an ERE may be significantly smaller than the DFA corresponding to the ERE (and obviously, it may have a different language).

Our next result shows that asynchronous monitoring of ERE-safety also

requires non-elementary space, which is a more general lower-bound result than the space non-elementarity of synchronous monitoring. Moreover, it gives an alternative proof of the non-elementary lower-bounds by Stockmeyer and Meyer [125, 124], because automata corresponding to safety-defining EREs are just special cases of monitors, so they must also take non-elementary space. Moreover, we improve the results in [125, 124] by showing that their lower-bounds also hold for a *subset of EREs*, namely those corresponding to safety properties.

Summarizing the discussion above, we believe that the main contributions of our subsequent lower-bound result are the following:

1. We show that asynchronous monitoring already requires non-elementary space, same as synchronous monitoring. For example, an ERE monitoring algorithm was presented by Roşu and Viswanathan in [103], which “rewrites” or “derives” the ERE by each letter in the input word; the derivation process consists of some straightforward rewrite rules, some for expanding the ERE others for contracting it via simplifications. No comprehensive and expensive check for emptiness on the resulting ERE is performed, except what is done by the simplification rules (for example  $\emptyset \cdot R \rightarrow \emptyset$  and  $\epsilon \cdot R \rightarrow R$ , etc.). A check for emptiness can and should be eventually performed (for example, a check for emptiness can be done periodically, say every  $10^x$  events for some convenient  $x$ ). This gives us an asynchronous ERE monitoring algorithm, which, unlike the simplistic NFA/DFA generation algorithm, does not pay upfront the potentially non-elementary worst-case penalty! However, our subsequent lower bound result tells that there is a worst-case scenario in which one cannot avoid the non-elementary space required to store the continuously changing ERE if one wants to correctly eventually detect violations of the original ERE. And that is the case for any synchronous or asynchronous monitoring algorithm for safety properties expressed as EREs.
2. We propose a different technique to prove non-elementary lower-bounds, fundamentally different from the one in [125]. The technique in [125] is based on diagonalization arguments and encodings of accepting Turing machine computations as finite trace words. Our technique is inspired from an idea by Chandra, Kozen and Stockmeyer [17] introduced to show the power of alternation and then used by several authors to prove exponential lower bounds [79, 80, 103, 105]. At our knowledge, the use of such a technique to show non-elementary lower bounds is

novel. The idea of the technique in [17] is to define expressions having as languages  $L_n = \{w^{(1)}\#w^{(2)}\#\dots\#w^{(k)}\$w \mid w^{(1)}, w^{(2)}, \dots, w^{(k)}, w \in \{0, 1\}^n, (\exists 1 \leq i \leq k) w^{(i)} = w\}$ . Our idea is to define, using EREs, words of the form  $X_n\$X'_n$ , where  $X_n$  and  $X'_n$  are  $n$ -deep nested sets starting with elements in  $\{0, 1\}^n$  (i.e., sets of sets of ... of sets of elements in  $\{0, 1\}^n$ , with  $n$  power set operations), such that  $X'_n$  is  $n$ -nested included in  $X_n$ , where  $i$ -nested inclusion is standard inclusion when  $i = 1$  and, if  $i > 1$ , then it is defined inductively as:  $X'_i$  is  $i$ -nested included in  $X_i$  iff for each  $X'_{i-1} \in X'_i$ , there is some  $X_{i-1} \in X_i$  such that  $X'_{i-1}$  is  $(i - 1)$ -nested included in  $X_{i-1}$ .

One more observation is in place before we move on to the technical details. It is known that the *membership problem* for EREs, testing whether a word  $w$  of size  $m$  is in the language of an ERE of size  $n$ , is polynomial in  $m$  and  $n$ . For example, the classic algorithm by Hopcroft and Ullman [70] runs in space  $O(m^2 \cdot n)$  and time  $O(m^3 \cdot n)$ ; slightly improved algorithms have been proposed by several authors [66, 130, 131, 132, 82, 73], reducing space requirements to  $O(m^2 \cdot k + m \cdot n)$  and time to  $O(m^3 \cdot k + m^2 \cdot n)$  or worse, where  $k$  is the number of complement operators in the ERE; a recent ERE membership algorithm was proposed by the author in [108], which runs in space  $O(m \cdot \log m \cdot 2^n)$  and time  $O(m^2 \cdot \log m \cdot 2^n)$  when  $m > 2^n$ . These algorithms appear to be efficient, because they are polynomial or simply exponential in the ERE, so one may think that one may device an elementary ERE-safety monitoring algorithm as follows: store the trace of events and at each newly received event invoke one of these “efficient” ERE membership algorithms. While this algorithm will indeed be elementary in the size of the ERE *and* the size of the trace, our lower bound result says that it will, in fact, be *non-elementary in the size of only the ERE!* In other words, for a carefully chosen “hard” ERE of size  $n$ , there are finite traces of large enough size  $m$  so that checking their membership is a problem which is non-elementary in  $n$ ; this will indeed happen when  $m$  is non-elementarily larger than  $n$ .

### 7.1.2 The Lower-Bound Result

EREs define languages by inductively applying *union* (+), *concatenation* ( $\cdot$ ), *Kleene Closure* ( $\star$ ), *intersection* ( $\cap$ ), and *complementation* ( $\neg$ ). The language of an ERE  $R$ , say  $\mathcal{L}(R)$ , is defined inductively as follows, where  $s \in \Sigma$ :

- $\mathcal{L}(\emptyset) = \emptyset$ ,
- $\mathcal{L}(\epsilon) = \{\epsilon\}$ ,
- $\mathcal{L}(s) = \{s\}$ ,
- $\mathcal{L}(R_1 + R_2) = \mathcal{L}(R_1) \cup \mathcal{L}(R_2)$ ,
- $\mathcal{L}(R_1 \cdot R_2) = \mathcal{L}(R_1) \cdot \mathcal{L}(R_2)$ ,
- $\mathcal{L}(R^*) = (\mathcal{L}(R))^*$ ,
- $\mathcal{L}(R_1 \cap R_2) = \mathcal{L}(R_1) \cap \mathcal{L}(R_2)$ ,
- $\mathcal{L}(\neg R) = \neg \mathcal{L}(R)$ .

If  $R$  does not contain  $\neg$  and  $\cap$  then it is a *regular expression* (RE). By applying De Morgan's law  $R_1 \cap R_2 \equiv \neg(\neg R_1 + \neg R_2)$ , EREs can be linearly (in both time and size) translated into equivalent EREs without intersection; therefore, intersection of EREs is just syntactic sugar. The *size* of an ERE is the total number of occurrences of letters and composition operators ( $+$ ,  $\cdot$ ,  $\star$ , and  $\neg$ ) that it contains. In what follows we assume that  $\Sigma$  is finite. For notational simplicity, in what follows we let  $\Sigma$  also denote the RE  $s_1 + s_2 + \dots + s_n$  where  $\Sigma = \{s_1, s_2, \dots, s_n\}$  and let  $\Sigma^*$  denote both the language  $\{s_1, s_2, \dots, s_n\}^*$  and the RE  $(s_1 + s_2 + \dots + s_n)^*$ .

For  $n \in \mathbb{N}$ , let us define inductively the following alphabets and languages:

- $\Sigma_0 = \{0, 1\}$  and  $\Psi_0 = \{0, 1\}^n$ , and
- $\Sigma_i = \Sigma_{i-1} \cup \{\#_i\}$  and  $\Psi_i = \{\#_i \#_i\} \cup (\{\#_i\} \cdot \Psi_{i-1})^+ \cdot \{\#_i\}$ , for all  $1 \leq i \leq n$ .

In the above,  $\#_i$  are  $n$  fresh letters. The intuition for the languages  $\Psi_i$  above is to encode nested sets of depth  $i$  that contain sets of words of  $n$  bits at their deepest level. The symbols  $\#_i$  play the role of markers separating the elements of such sets. For example, the word  $\#_2 \#_1 \#_1 \#_2 \#_1 01 \#_1 10 \#_1 \#_2 \#_1 \#_1 \#_2$  encodes  $\{\{\}, \{01, 10\}, \{\}\}$ , that is, the set  $\{\{\}, \{01, 10\}\}$ ; since the multiplicity and order of elements in sets are irrelevant, the same set can have (infinitely) many different encodings. Formally, let us define the following *set* functions associating to encodings in  $\Psi_i$  corresponding nested sets:

- $set_0 : \Psi_0 \rightarrow \{0, 1\}^n$  is the identity function on  $\Psi_0$ , i.e.,  $set_0(w) = w$ ;

- $set_i : \Psi_i \rightarrow \mathcal{P}^i(\{0, 1\}^n)$ , where  $\mathcal{P}^i$  is the power set operator applied  $i$  times,  $set_i(\#_i \#_i) = \{\}$ ,  $set_i(\#_i X_{i-1} \#_i) = \{set_{i-1}(X_{i-1})\}$ , and  $set_i(\#_i X_{i-1} X_i) = \{set_{i-1}(X_{i-1})\} \cup set_i(X_i)$ , for all  $1 \leq i \leq n$ ,  $X_{i-1} \in \Psi_{i-1}$ , and  $X_i \in \Psi_i$ .

Note that  $|set_0(\Psi_0)| = 2^n$  and  $set_i(\Psi_i) = \mathcal{P}(set_{i-1}(\Psi_{i-1}))$  for all  $1 \leq i \leq n$ ; therefore,  $|set_i(\Psi_i)| = 2^{2^{i-1} \cdot 2^n}$  for all  $1 \leq i \leq n$ , with  $i + 1$  nested power operations.

Let us define *nested-inclusion* relations  $\_ \subseteq_i \_ : \mathcal{P}^i(\{0, 1\}^n) \times \mathcal{P}^i(\{0, 1\}^n)$  for  $0 \leq i \leq n$  and *nested-membership* relations  $\_ \in_i \_ : \mathcal{P}^{i-1}(\{0, 1\}^n) \times \mathcal{P}^i(\{0, 1\}^n)$  for  $1 \leq i \leq n$  as follows:

- $\_ \subseteq_0 \_$  is the identity on  $\{0, 1\}^n$  and  $\_ \subseteq_1 \_$  is  $\_ \subseteq \_ : \mathcal{P}(\{0, 1\}^n) \times \mathcal{P}(\{0, 1\}^n)$ ,
- $\_ \in_1 \_$  is  $\_ \in \_ : \{0, 1\}^n \times \mathcal{P}(\{0, 1\}^n)$ , and for  $1 < i$ ,
- $S_{i-1} \in S_i$  iff there is some  $S'_{i-1} \in S_i$  such that  $S_{i-1} \subseteq S'_{i-1}$ , and
- $S_i \subseteq S'_i$  iff  $S_{i-1} \in S'_i$  for each  $S_{i-1} \in S_i$ .

For example, if  $n = 2$  then  $\{\{00, 01\}, \{01, 10\}, \{11\}\} \subseteq_2 \{\{00, 01, 10\}, \{00, 11\}\}$  because  $\{00, 01\}, \{01, 10\} \subseteq_1 \{00, 01, 10\}$  and  $\{11\} \subseteq_1 \{00, 11\}$ .

We can now define  $\Sigma$  as  $\Sigma_n \cup \{\$ \}$  and the infinite trace property  $P_n^\omega$  over  $\Sigma$ :

$$(\epsilon \cup (\Sigma_n^* \cup \{X_n \$ X'_n \mid X_n, X'_n \in \Psi_n \text{ and } set_n(X'_n) \subseteq_n set_n(X_n)\}) \cdot \{\$ \}) \cdot \Sigma_n^\omega.$$

An infinite trace in  $P_n^\omega$  therefore contains at most two  $\$$  letters and infinitely many letters in  $\Sigma_n$ . There are no restrictions on the appearance of the letters in  $\Sigma_n$  when there is no  $\$$  letter or when there is precisely one  $\$$  letter. However, if the infinite trace contains precisely two  $\$$  letters, that is, if it has the form  $w \$ w' \$ u$  for some  $w, w' \in \Sigma_n^*$  and some  $u \in \Sigma_n^\omega$ , then  $w$  and  $w'$  must be in  $\Psi_n$  and the nested set corresponding to  $w$  must nested-include the nested set corresponding to  $w'$ ; there are no restrictions on  $u$ .

**Proposition 15**  $P_n^\omega \in \text{Safety}^\omega$ .

**Proof:** There are two cases to analyze for an infinite trace that is not in  $P_n^\omega$ : when it contains more than two  $\$$  letter, or when it has the form

$w\$w'\$u$  with  $w, w' \in \Sigma_n^*$  and  $u \in \Sigma_n^\omega$ , but it is not the case that  $w, w' \in \Psi_n$  and  $set_n(w') \subseteq_n set_n(w)$ . In the first case, we can pick the first prefix of the infinite trace containing three \$ letters in total; clearly, this finite trace prefix cannot be continued into any acceptable infinite trace. In the second case, since there are no restrictions on the letters in  $u$ , we can easily see that the prefix  $w\$w'\$$  is already a violation threshold: there is no  $u' \in \Sigma_n^\omega$  such that  $w\$w'\$u' \in P_n^\omega$ .  $\square$

The bijection in the proof of Theorem 3 associates to each infinite-trace safety property a persistent finite-trace safety property by taking its prefixes. Let  $P_n$  be the persistent finite-trace safety property  $\text{prefixes}(P_n^\omega)$  corresponding to  $P_n^\omega$ . It is easy to see that  $P_n$  is the property

$$\Sigma_n^* \cup \Sigma_n^* \cdot \{\$ \} \cdot \Sigma_n^* \cup \{X_n \$ X'_n \mid X_n, X'_n \in \Psi_n \text{ and } set_n(X'_n) \subseteq_n set_n(X_n)\} \cdot \{\$ \} \cdot \Sigma_n^*.$$

Note that monitoring  $P_n^\omega$  is the same as monitoring  $P_n$ : in both cases, besides the capability to checking whether there are more than two \$ letters, which is trivial, the monitor needs to store sufficient information about the nested set corresponding to  $X_n$ , so that, when the first \$ is seen, to be able to check whether it nested-contains the set corresponding to the upcoming, yet unknown  $X'_n$ .

**Theorem 8** *Any synchronous or asynchronous monitor for  $P_n$  or  $P_n^\omega$  needs space non-elementary in  $n$ , namely  $\Omega(2^{2^{2^n}})$ , with  $n$  nested power operations.*

**Proof:** Suppose that  $M$  is a monitor for  $P_n$  or  $P_n^\omega$  and suppose that, during a monitoring session, it reads the prefix  $X_n \in \Psi_n$ . Regardless of how  $M$  is defined or implemented, in particular regardless of whether it reports violations synchronously or asynchronously, when the first \$ event is encountered, the state of  $M$  must contain enough information to sooner or later be able to decide whether the set  $set_n(X'_n)$  corresponding to *any* upcoming (unknown at the time the \$ is observed) sequence  $X'_n$  is nested-contained in  $set_n(X_n)$ . Since  $set(X'_n)$  can in particular be equal to  $set_n(X_n)$ , and since once the second \$ event is observed there is no further event that may bring new knowledge to the monitor, we deduce that  $M$  must be able to distinguish any two different sets in  $set_n(\Psi_n)$  when the first \$ event is encountered, that is,  $M$ 's states must be different after reading words in  $\Psi_n$  whose corresponding nested sets are different. Therefore,  $M$  must be able to distinguish  $|set_n(\Psi_n)|$  different possibilities. Since one needs  $\Omega(\log N)$

space to distinguish among  $N$  different situations (one label for each), we conclude that  $M$  needs space  $\Omega(\log(|\text{set}_n(\Psi_n)|))$ , that is,  $\Omega(2^{2^{2^n}})$  with  $n$  nested power operations. Hence, any monitor for  $P_n$  needs non-elementarily large space in  $n$ .  $\square$

We next show how to construct an ERE polynomial in size with  $n$  whose language is precisely  $P_n$ .

**Theorem 9** *There is an ERE of size  $O(n^3)$  whose language is  $P_n$ .*

**Proof:** Note that  $P_n$  is a union of three languages, the first two being trivial to express as languages of corresponding REs. As expected, the difficult part is to associate an ERE to the language

$$\{X_n \$ X'_n \mid X_n, X'_n \in \Psi_n \text{ and } \text{set}_n(X'_n) \subseteq_n \text{set}_n(X_n)\}.$$

Note that so far we did not need complementation. The property above can, however, be expressed as an ERE of size  $O(n^3)$  using  $O(n)$  nested complement operators. The idea is to define iteratively a sequence of EREs  $\mathbb{K}_i$  for  $0 \leq i \leq n$  whose languages contain words  $X_i w \$ w' X'_i$  with  $\text{set}_i(X'_i) \subseteq_i \text{set}(X_i)$ , which are contiguous fragments of desired words  $X_n \$ X'_n$ . Then  $\mathbb{K}_n$  would be the language that we are looking for. To define  $\mathbb{K}_i$ , we observe that the nested-inclusion  $S'_i \subseteq_i S_i$  is equivalent to: there is no  $S'_{i-1} \in S'_i$  such that it is not the case that we can find some  $S_{i-1} \in S_i$  such that  $S'_{i-1} \subseteq_{i-1} S_{i-1}$ . This crucial observation will allow us to define  $\mathbb{K}_i$  in terms of  $\mathbb{K}_{i-1}$ . We next develop the technical details.

Let us first define regular patterns corresponding to each of the languages  $\Psi_i$  for  $0 \leq i \leq n$ ; to avoid introducing new names, we ambiguously let the corresponding regular expressions have the same names as their languages:

- Let  $\Psi_0 = (0 + 1)^n$ , where for an RE,  $r^n$  is  $r \cdot r \cdot \dots \cdot r$  ( $n$  times); and
- Let  $\Psi_i = \#_i \cdot \#_i + (\#_i \cdot \Psi_{i-1})^+ \cdot \#_i$  for all  $1 \leq i \leq n$ .

Iteratively eliminating the  $\Psi_{i-1}$  regular expressions from the right-hand-sides, we eventually obtain  $n + 1$  regular patters, each of size  $O(n)$  (the size of  $\Psi_0$  as a regular expression is  $O(n)$  and each  $\Psi_i$  adds a constant size to that of  $\Psi_{i-1}$ ).

Next we define REs for the languages  $\text{prefixes}(\Psi_i)$  and  $\text{suffixes}(\Psi_i)$  of prefixes and respectively suffixes of words in  $\Psi_i$ , for all  $0 \leq i \leq n$ . We only discuss the prefix closure languages; the suffix closures are dual. The prefix closures can be defined relatively easily inductively as follows:

- $\text{prefixes}(\Psi_0) = \bigcup_{k=0}^n \{0, 1\}^k$ , and
- $\text{prefixes}(\Psi_i) = \{\epsilon, \#_i \#_i\} \cup \{\#_i\} \cdot \text{prefixes}(\Psi_{i-1}) \cup (\{\#_i\} \cdot \Psi_{i-1})^+ \cup (\{\#_i\} \cdot \Psi_{i-1})^+ \cdot \{\#_i\} \cdot \text{prefixes}(\Psi_{i-1})$   
 $= \{\#_i \#_i\} \cup (\{\#_i\} \cdot \Psi_{i-1})^* \cdot (\{\epsilon\} \cup \{\#_i\} \cdot \text{prefixes}(\Psi_{i-1}))$ .

These languages can be expressed with the following REs; as before, we ambiguously use the same names for the corresponding REs:

- Let  $\text{prefixes}(\Psi_0) = \epsilon + (0 + 1) + (0 + 1)^2 + \dots + (0 + 1)^n = (\epsilon + 0 + 1)^n$ ; and
- Let  $\text{prefixes}(\Psi_i) = \#_i \cdot \#_i + (\#_i \cdot \Psi_{i-1})^* \cdot (\epsilon + \#_i \cdot \text{prefixes}(\Psi_{i-1}))$ .

Iteratively eliminating the REs  $\text{prefixes}(\Psi_{i-1})$  from the right-hand-sides, we eventually obtain  $n + 1$  REs, each of size  $O(i^2 + n)$  (the size of  $\text{prefixes}(\Psi_0)$  as an RE is  $O(n)$  and each  $\text{prefixes}(\Psi_i)$  adds size  $O(i)$  to that of  $\text{prefixes}(\Psi_{i-1})$ ). Dually,

- Let  $\text{suffixes}(\Psi_0) = \epsilon + (0 + 1) + (0 + 1)^2 + \dots + (0 + 1)^n = (\epsilon + 0 + 1)^n$ ; and
- Let  $\text{suffixes}(\Psi_i) = \#_i \cdot \#_i + (\epsilon + \text{suffixes}(\Psi_{i-1}) \cdot \#_i) \cdot (\Psi_{i-1} \cdot \#_i)^*$ .

We next define REs  $L_i$  and  $R_i$  for  $0 \leq i \leq n$  whose languages contain the contiguous fragments of words in  $\Psi_n$  that are allowed to appear to the left and to the right of  $\$,$  respectively, so that words in  $\mathcal{L}(L_i)$  start with  $\#_i$  and words in  $\mathcal{L}(R_i)$  end with  $\#_i$ . Let us also assume by convention that  $L_{n+1} = R_{n+1} = \epsilon$  (the RE whose language contains only the empty word). It is easy to see that  $L_i$  and  $R_i$  can be defined as follows:

- Let  $L_i = \#_i \cdot \Sigma_n^* \cap \text{suffixes}(\Psi_n)$ , and
- Let  $R_i = \Sigma_n^* \cdot \#_i \cap \text{prefixes}(\Psi_n)$ .

Note that words in  $L_i$  and  $R_i$  may not necessarily start or end with a word in  $\Psi_i$ : indeed, the  $\#_i$  that may start or end  $L_i$  or  $R_i$  could very well be followed or preceded, respectively, by a  $\#_{i+1}$  or, if  $i = n$ , by  $\$$ .

Let us also define the EREs  $\bar{L}_i$  and  $\bar{R}_i$  whose languages are included in those of  $L_i$  and  $R_i$ , respectively, and whose words start or end with words in  $\Psi_i$ :

- Let  $\bar{L}_i = \Psi_i \cdot \Sigma_n^* \cap \text{suffixes}(\Psi_n)$ , and

- Let  $\bar{R}_i = \Sigma_n^* \cdot \Psi_i \cap \text{prefixes}(\Psi_n)$ .

It is not difficult to see that  $\bar{L}_i = \Psi_i \cdot L_{i+1}$  and  $\bar{R}_i = R_{i+1} \cdot \Psi_i$ . Note that the sizes of  $L_i$ ,  $R_i$ ,  $\bar{L}_i$  and  $\bar{R}_i$  are  $O(n^2)$ .

Let us now define the EREs  $\mathbb{K}_i$  for  $0 \leq i \leq n$  as follows:

- $\mathbb{K}_0 = \bar{L}_0 \cdot \$ \cdot \bar{R}_0 \cap \bigcap_{k=0}^{n-1} (\Sigma_0^k \cdot 0 \cdot \Sigma^* \cdot 0 \cdot \Sigma_0^{n-k-1} + \Sigma_0^k \cdot 1 \cdot \Sigma^* \cdot 1 \cdot \Sigma_0^{n-k-1})$ ,  
and
- $\mathbb{K}_i = \bar{L}_i \cdot \$ \cdot \bar{R}_i \cap \neg((\neg((\#_i \cdot \Psi_{i-1})^* \cdot \#_i \cdot \mathbb{K}_{i-1}) \cap L_i \cdot \$ \cdot R_{i-1}) \cdot (\#_i \cdot \Psi_{i-1})^* \cdot \#_i)$ .

We next show by induction on  $i$  that  $\mathcal{L}(\mathbb{K}_i)$  is the language

$$\{X_i w X'_i \mid X_i, X'_i \in \Psi_i, w \in \mathcal{L}(L_{i+1} \cdot \$ \cdot R_{i+1}), \text{set}_i(X'_i) \subseteq_i \text{set}_i(X_i)\}.$$

It is easy to see that  $\mathcal{L}(\mathbb{K}_0) = \{X_0 w X_0 \mid X_0 \in \Psi_0, w \in \mathcal{L}(L_1 \cdot \$ \cdot R_1)$ , because the large conjunct in  $\mathbb{K}_0$  states that the words formed with the first  $n$  letters and with the last ones, respectively, are equal and in  $\Psi_0$ , and because  $\bar{L}_0 \cdot \$ \cdot \bar{R}_0 = \Psi_0 \cdot L_1 \cdot \$ \cdot R_1 \cdot \Psi_0$  and  $\subseteq_0$  is the identity on  $\{0, 1\}^n$ . For the inductive step, let us now assume that for some arbitrary  $1 \leq i < n$ ,  $\mathcal{L}(\mathbb{K}_{i-1})$  is the language

$$\{X_{i-1} w X'_{i-1} \mid X_{i-1}, X'_{i-1} \in \Psi_{i-1}, w \in \mathcal{L}(L_i \cdot \$ \cdot R_i), \text{set}_{i-1}(X'_{i-1}) \subseteq_{i-1} \text{set}_{i-1}(X_{i-1})\}.$$

Then we can easily show that  $\mathcal{L}((\#_i \cdot \Psi_{i-1})^* \cdot \#_i \cdot \mathbb{K}_{i-1})$  is the language

$$\{X_i w X'_{i-1} \mid X_i \in \Psi_i, X'_{i-1} \in \Psi_{i-1}, w \in \mathcal{L}(L_{i+1} \cdot \$ \cdot R_i), \text{set}_{i-1}(X'_{i-1}) \subseteq_{i-1} \text{set}_i(X_i)\}.$$

Then we can show that  $\mathcal{L}(\neg((\#_i \cdot \Psi_{i-1})^* \cdot \#_i \cdot \mathbb{K}_{i-1}) \cap \bar{L}_i \cdot \$ \cdot \bar{R}_{i-1})$  is

$$\{X_i w X'_{i-1} \mid X_i \in \Psi_i, X'_{i-1} \in \Psi_{i-1}, w \in \mathcal{L}(L_{i+1} \cdot \$ \cdot R_i), \text{set}_{i-1}(X'_{i-1}) \not\subseteq_{i-1} \text{set}_i(X_i)\}.$$

Now we can show that  $\mathcal{L}(\neg((\#_i \cdot \Psi_{i-1})^* \cdot \#_i \cdot \mathbb{K}_{i-1}) \cap \bar{L}_i \cdot \$ \cdot \bar{R}_{i-1}) \cdot (\#_i \cdot \Psi_{i-1})^* \cdot \#_i$  is the language

$$\{X_i w X'_i \mid X_i, X'_i \in \Psi_i, w \in \mathcal{L}(L_{i+1} \cdot \$ \cdot R_{i+1}), \text{set}_i(X'_i) \not\subseteq_{i-1} \text{set}_i(X_i)\}.$$

Finally, we are now able to show that  $\mathcal{L}(\mathbb{K}_i)$ , that is,

$$\mathcal{L}(\bar{L}_i \cdot \$ \cdot \bar{R}_i \cap \neg((\neg((\#_i \cdot \Psi_{i-1})^* \cdot \#_i \cdot \mathbb{K}_{i-1}) \cap L_i \cdot \$ \cdot R_{i-1}) \cdot (\#_i \cdot \Psi_{i-1})^* \cdot \#_i))$$

is the language

$$\{X_i w X'_i \mid X_i, X'_i \in \Psi_i, w \in \mathcal{L}(L_{i+1} \cdot \$ \cdot R_{i+1}), \text{set}_i(X'_i) \subseteq_i \text{set}_i(X_i)\}.$$

Since  $L_{n+1} = R_{n+1} = \epsilon$ , it follows that

$$\mathcal{L}(\mathbb{K}_n) = \{X_n \$ X'_n \mid X_n, X'_n \in \Psi_n, set_i(X'_i) \sqsubseteq_i set_i(X_i)\}.$$

The size of  $\mathbb{K}_n$  is  $O(n^3)$ .

We can now show that the language of the ERE of size  $O(n^3)$

$$(\epsilon + (\Sigma_n^* + \mathbb{K}_n) \cdot \$) \cdot \Sigma_n^*$$

is indeed  $P_n$ .

We can now formulate our space lower-bound result for monitoring ERE-safety as a corollary of the two results above.

**Corollary 1** *For any  $n \in \mathbb{N}$ , there is some safety property whose good prefixes are precisely the words in the language of an ERE of size  $O(n)$  and whose monitoring (synchronous or asynchronous) requires space  $\Omega(2^{2^{\cdot_2 \sqrt[3]{n}}})$  with  $\sqrt[3]{n}$  nested power operations.*

## 7.2 Generating Optimal Monitors for ERE

*incorporate RTA '03 paper with Mahesh*

---

*Abstract: Software engineers and programmers can easily understand regular patterns, as shown by the immense interest in and the success of scripting languages like Perl, based essentially on regular expression pattern matching. We believe that regular expressions provide an elegant and powerful specification language also for monitoring requirements, because an execution trace of a program is in fact a string of states. Extended regular expressions (EREs) add complementation to regular expressions, which brings additional benefits by allowing one to specify patterns that must not occur during an execution. Complementation gives one the power to express patterns on strings more compactly. In this paper we present a technique to generate optimal monitors from EREs. Our monitors are deterministic finite automata (DFA) and our novel contribution is to generate them using a modern coalgebraic technique called coinduction. Based on experiments with our implementation, which can be publicly tested and used over the web, we believe that our technique is more efficient than the simplistic method based on complementation of automata which can quickly lead to a highly-exponential state explosion.*

---

### 7.2.1 Introduction

Regular expressions can express patterns in strings in a compact way. They proved very useful in practice; many programming/scripting languages like Perl, Python, Tcl/Tk support regular expressions as core features. Because of their power to express a rich class of patterns, regular expressions, are used not only in computer science but also in various other fields, such as molecular biology [77]. All these applications boast of very efficient implementation of regular expression pattern matching and/or membership algorithms. Moreover, it has been found that compactness of regular expressions can be increased non-elementarily by adding complementation ( $\neg R$ ) to the usual union ( $R_1 + R_2$ ), concatenation ( $R_1 \cdot R_2$ ), and repetition ( $R^*$ ) operators of regular expressions. These are known as *extended regular expressions* (EREs) and they proved very intuitive and succinct in expressing regular patterns.

Recent trends have shown that the software analysis community is inclining towards scalable techniques for software verification. Works in [58] merged temporal logics with testing, hereby getting the benefits of both

worlds. The Temporal Rover tool (TR) and its follower DB Rover [32] are already commercial. In these tools the Java code is instrumented automatically so that it can check the satisfaction of temporal logic properties at runtime. The MaC tool [76, 87] has been developed to monitor safety properties in interval past time temporal logics. In [98, 101], various algorithms to generate testing automata from temporal logic formulae, are described. Java PathExplorer [55] is a runtime verification environment currently under development at NASA Ames. The Java MultiPathExplorer tool [116] proposes a technique to monitor all equivalent traces that can be extracted from a given execution, thus increasing the coverage of monitoring. [44, 57] present efficient algorithms for monitoring future time temporal logic formulae, while [60] gives a technique to synthesize efficient monitors from past time temporal formulae. [103] uses rewriting to perform runtime monitoring of EREs.

An interesting aspect of EREs is that they can express safety properties compactly, like those encountered in testing and monitoring. By generating automata from logical formulae, several of the works above show that the safety properties expressed by different variants of temporal logics are subclasses of regular languages. The converse is *not* true, because there are regular patterns which cannot be expressed using temporal logics, most notoriously those related to counting; e.g., the regular expression  $(0 \cdot (0+1))^*$  saying that every other letter is 0 does not admit an equivalent temporal logic formula. Additionally, EREs tend to be often very natural and intuitive in expressing requirements. For example, let us try to capture the safety property “it should not be the case that in any trace of a traffic light we see green and then immediately red at any point”. The natural and intuitive way to express it in ERE is  $\neg((-\emptyset) \cdot \text{green} \cdot \text{red} \cdot (-\emptyset))$ , where  $\emptyset$  is the empty ERE (no words), so  $-\emptyset$  means “anything”.

Previous approaches to ERE membership testing [66, 96, 130, 82, 74] have focussed on developing techniques that are polynomial in both the size of the word and the size of the formulae. The best known result in these approaches is described in [82] where they can check if a word satisfies an ERE in time  $O(m \cdot n^2)$  and space  $O(m \cdot m + k \cdot n^2)$ , where  $m$  is the size of the ERE,  $n$  is the length of the word, and  $k$  is the number of negation/intersection operators. These algorithms, unfortunately, cannot be used for the purpose of monitoring. This is because they are not incremental. They assume the entire word is available before their execution. Additionally, their running time and space requirements are quadratic in the size of the trace. This

is unacceptable when one has a long trace of events and wants to monitor a small ERE, as it is typically the case. This problem is removed in [103] where traces are checked against EREs through incremental rewriting. At present, we do not know if the technique in [103] is optimal or not.

A simple, straightforward, and practical approach is to generate optimal *deterministic finite automata* (DFA) from EREs [69]. This process involves the conversion of each negative sub-component of the ERE to a non-deterministic finite automaton (NFA), determinization of the NFA into a DFA, complementation of the DFA, and then its minimization. The algorithm runs in a bottom-up fashion starting from the innermost negative ERE sub components. This method, although generates the minimal automata, is too complex and cumbersome in practice. Its space requirements can be non-elementarily larger than the initial regular ERE, because negation involves an NFA-to-DFA translation, which implies an exponential blow-up; since negations can be nested, the size of such NFAs or DFAs could be highly exponential.

Our approach is to generate the minimal DFA from an ERE using coinductive techniques. In this paper, the DFA thus generated is called the *optimal monitor* for the given ERE. Currently, we are not aware of any other algorithm that does this conversion in a straightforward way. The complexity of our algorithm seems to be hard to evaluate, because it depends on the size of the minimal DFA associated to an ERE and we are not aware of any lower bound results in this direction. However, experiments are very encouraging. Our implementation, which is available for evaluation on the internet via a CGI server reachable from <http://fs1.cs.uiuc.edu/rv/>, rarely took longer than one second to generate a DFA, and it took only 18 minutes to generate the minimal 107 state DFA for the ERE in Example 13 which was used to show the exponential space lower bound of ERE monitoring in [103].

In a nutshell, in our approach we use the concept of derivatives of an ERE, as described in Subsection 7.2.2. For a given ERE one generates all possible derivatives of the ERE for all possible sequences of events. The size of this set of derivatives depends upon the size of the initial ERE. However, several of these derivative EREs can be equivalent to each other. One can check the equivalence of EREs using coinductive technique as described in Section 7.2.3, that generates a set of equivalent EREs, called *circularities*. In Section 7.2.5, we show how circularities can be used to construct an efficient algorithm that generates optimal DFAs from EREs. In Section 7.2.6, we describe an implementation of this algorithm and give performance analysis

results. We also made available on the internet a CGI interface to this algorithm.

### 7.2.2 Extended Regular Expressions and Derivatives

In this section we recall extended regular expressions and their derivatives.

#### Extended Regular Expressions

Extended regular expressions (ERE) define languages by inductively applying union (+), concatenation ( $\cdot$ ), Kleene Closure (\*), intersection ( $\cap$ ), and complementation ( $\neg$ ). More precisely, for an alphabet  $E$ , whose elements are called *events* in this paper, an ERE over  $E$  is defined as follows, where  $a \in E$ :

$$R ::= \emptyset \mid \epsilon \mid a \mid R + R \mid R \cdot R \mid R^* \mid R \cap R \mid \neg R.$$

The language defined by an expression  $R$ , denoted by  $\mathcal{L}(R)$ , is defined inductively as

$$\begin{aligned} \mathcal{L}(\emptyset) &= \emptyset, \\ \mathcal{L}(\epsilon) &= \{\epsilon\}, \\ \mathcal{L}(A) &= \{A\}, \\ \mathcal{L}(R_1 + R_2) &= \mathcal{L}(R_1) \cup \mathcal{L}(R_2), \\ \mathcal{L}(R_1 \cdot R_2) &= \{w_1 \cdot w_2 \mid w_1 \in \mathcal{L}(R_1) \text{ and } w_2 \in \mathcal{L}(R_2)\}, \\ \mathcal{L}(R^*) &= (\mathcal{L}(R))^*, \\ \mathcal{L}(R_1 \cap R_2) &= \mathcal{L}(R_1) \cap \mathcal{L}(R_2), \\ \mathcal{L}(\neg R) &= \Sigma^* \setminus \mathcal{L}(R). \end{aligned}$$

Given an ERE, as defined above using union, concatenation, Kleene Closure, intersection and complementation, one can translate it into an equivalent expression that does not have any intersection operation, by applying De Morgan's Law:  $R_1 \cap R_2 = \neg(\neg R_1 + \neg R_2)$ . The translation only results in a linear blowup in size. Therefore, in the rest of the paper we do not consider expressions containing intersection. More precisely, we only consider EREs of the form

$$R ::= R + R \mid R \cdot R \mid R^* \mid \neg R \mid a \mid \epsilon \mid \emptyset.$$

*Derivatives*

In this subsection we recall the notion of *derivative*, or “residual” (see [7, 6], where several interesting properties of derivatives are also presented). It is based on the idea of “event consumption”, in the sense that an extended regular expression  $R$  and an event  $a$  produce another extended regular expression, denoted  $R\{a\}$ , with the property that for any trace  $w$ ,  $aw \in R$  if and only if  $w \in R\{a\}$ .

In the rest of the paper assume defined the typical operators on EREs and consider that the operator  $_ + _$  is associative and commutative and that the operator  $_ \cdot _$  is associative. In other words, reasoning is performed modulo the equations:

$$\begin{aligned} (R_1 + R_2) + R_3 &= R_1 + (R_2 + R_3), \\ R_1 + R_2 &= R_2 + R_1, \\ (R_1 \cdot R_2) \cdot R_3 &= R_1 \cdot (R_2 \cdot R_3). \end{aligned}$$

We next consider an operation  $_ \{ \_ \}$  which takes an ERE and an event, and give several equations which define its operational semantics recursively, on the structure of regular expressions:

$$\begin{aligned} (R_1 + R_2)\{a\} &= R_1\{a\} + R_2\{a\} & (1) \\ (R_1 \cdot R_2)\{a\} &= (R_1\{a\}) \cdot R_2 + \text{if } (\epsilon \in R_1) \text{ then } R_2\{a\} \text{ else } \emptyset \text{ fi} & (2) \\ (R^*)\{a\} &= (R\{a\}) \cdot R^* & (3) \\ (\neg R)\{a\} &= \neg(R\{a\}) & (4) \\ b\{a\} &= \text{if } (b == a) \text{ then } \epsilon \text{ else } \emptyset \text{ fi} & (5) \\ \epsilon\{a\} &= \emptyset & (6) \\ \emptyset\{a\} &= \emptyset & (7) \end{aligned}$$

The right-hand sides of these equations use operations which we describe next. “if  $(\_)$  then  $\_$  else  $\_$  fi” takes a boolean term and two EREs as arguments and has the expected meaning defined by two equations:

$$\begin{aligned} \text{if } (true) \text{ then } R_1 \text{ else } R_2 \text{ fi} &= R_1 & (8) \\ \text{if } (false) \text{ then } R_1 \text{ else } R_2 \text{ fi} &= R_2 & (9) \end{aligned}$$

We assume a set of equations that properly define boolean expressions and reasoning. Boolean expressions include the constants *true* and *false*, as well as the usual connectors  $_ \wedge _$ ,  $_ \vee _$ , and *not*. Testing for empty trace membership (which is used by (2)) can be defined via the following equations:

$$\epsilon \in (R_1 + R_2) = (\epsilon \in R_1) \vee (\epsilon \in R_2) \quad (10)$$

$$\epsilon \in (R_1 \cdot R_2) = (\epsilon \in R_1) \wedge (\epsilon \in R_2) \quad (11)$$

$$\epsilon \in (R^*) = \text{true} \quad (12)$$

$$\epsilon \in (\neg R) = \text{not}(\epsilon \in R) \quad (13)$$

$$\epsilon \in b = \text{false} \quad (14)$$

$$\epsilon \in \epsilon = \text{true} \quad (15)$$

$$\epsilon \in \emptyset = \text{false} \quad (16)$$

The 16 equations above are natural and intuitive. [103] shows that these equations, when regarded as rewriting rules are terminating and ground Church-Rosser (modulo associativity and commutativity of  $_+ _$  and modulo associativity of  $_\cdot _$ ), so they can be used as a functional procedure to calculate derivatives. Due to the fact that the 16 equations defining the derivatives can generate useless terms, in order to keep EREs compact we also propose defining several *simplifying equations*, including at least the following:

$$\emptyset + R = R,$$

$$\emptyset \cdot R = \emptyset,$$

$$\epsilon \cdot R = R,$$

$$R + R = R.$$

The following result (see, e.g., [103] for a proof) gives a simple procedure, based on derivatives, to test whether a word belongs to the language of an ERE:

**Theorem 10** *For any ERE  $\mathcal{R}$  and any events  $a, a_1, a_2, \dots, a_n$  in  $A$ , the following hold:*

1.  $a_1 a_2 \dots a_n \in \mathcal{L}(R\{a\})$  if and only if  $aa_1 a_2 \dots a_n \in \mathcal{L}(R)$ ; and
2.  $a_1 a_2 \dots a_n \in \mathcal{L}(R)$  if and only if  $\epsilon \in R\{a_1\}\{a_2\}\dots\{a_n\}$ .

### 7.2.3 Hidden Logic and Coinduction

We use circular coinduction, defined rigorously in the context of hidden logics and implemented in the BOBJ system [102, 46, 47], to test whether two EREs are equivalent, that is, if they have the same language. Since the goal of this paper is to translate an ERE into a minimal DFA, standard techniques for checking equivalence, such as translating the two expressions into DFAs and then comparing those, do not make sense in this framework. A particularly appealing aspect of circular coinduction in the framework of

EREs is that it does not only show that two EREs are equivalent, but also generates a larger set of equivalent EREs which will all be used in order to generate the target DFA.

Hidden logic is a natural extension of algebraic specification which benefits of a series of generalizations in order to capture various natural notions of behavioral equivalence found in the literature. It distinguishes *visible* sorts for data from *hidden* sorts for states, with states *behaviorally equivalent* if and only if they are indistinguishable under a formally given set of experiments. To keep the presentation simple and self contained, in this section we define an oversimplified version of hidden logic together with its associated circular coinduction proof rule, still general enough to support defining and proving EREs behaviorally equivalent.

#### *Algebraic Preliminaries*

The reader is assumed familiar with basic equational logic and algebra in this section. We recall a few notions in order to just make our notational conventions precise. An  $S$ -sorted signature  $\Sigma$  is a set of sorts/types  $S$  together with operational symbols on those, and a  $\Sigma$ -algebra  $A$  is a collection of sets  $\{A_s \mid s \in S\}$  and a collection of functions appropriately defined on those sets, one for each operational symbol. Given an  $S$ -sorted signature  $\Sigma$  and an  $S$ -indexed set of variables  $Z$ , let  $T_\Sigma(Z)$  denote the  $\Sigma$ -term algebra over variables in  $Z$ . If  $V \subseteq S$  then  $\Sigma|_V$  is a  $V$ -sorted signature consisting of all those operations in  $\Sigma$  with sorts entirely in  $V$ . We may let  $\sigma(X)$  denote the term  $\sigma(x_1, \dots, x_n)$  when the number of arguments of  $\sigma$  and their order and sorts are not important. If only one argument is important, then to simplify writing we place it at the beginning; for example,  $\sigma(t, X)$  is a term having  $\sigma$  as root with only variables as arguments except one, and we do not care which one, which is  $t$ . If  $t$  is a  $\Sigma$ -term of sort  $s'$  over a special variable  $*$  of sort  $s$  and  $A$  is a  $\Sigma$ -algebra, then  $A_t : A_s \rightarrow A_{s'}$  is the usual interpretation of  $t$  in  $A$ .

#### **7.2.4 Behavioral Equivalence, Satisfaction and Specification**

Given disjoint sets  $V, H$  called *visible* and *hidden sorts*, a *hidden*  $(V, H)$ -signature, say  $\Sigma$ , is a many sorted  $(V \cup H)$ -signature. A *hidden subsignature* of  $\Sigma$  is a hidden  $(V, H)$ -signature  $\Gamma$  with  $\Gamma \subseteq \Sigma$  and  $\Gamma|_V = \Sigma|_V$ . The *data signature* is  $\Sigma|_V$ . An operation of visible result not in  $\Sigma|_V$  is called an *attribute*, and a hidden sorted operation is called a *method*.

Unless otherwise stated, the rest of this section assumes fixed a hidden signature  $\Sigma$  with a fixed subsignature  $\Gamma$ . Informally,  $\Sigma$ -algebras are universes

of possible states of a system, i.e., “black boxes,” where one is only concerned with behavior under experiments with operations in  $\Gamma$ , where an experiment is an observation of a system attribute after perturbation; this is formalized below.

A  $\Gamma$ -context for sort  $s \in V \cup H$  is a term in  $T_\Gamma(\{ * : s \})$  with one occurrence of  $*$ . A  $\Gamma$ -context of visible result sort is called a  $\Gamma$ -experiment. If  $c$  is a context for sort  $h$  and  $t \in T_{\Sigma, h}$  then  $c[t]$  denotes the term obtained from  $c$  by substituting  $t$  for  $*$ ; we may also write  $c[*]$  for the context itself.

Given a hidden  $\Sigma$ -algebra  $A$  with a hidden subsignature  $\Gamma$ , for sorts  $s \in (V \cup H)$ , we define  $\Gamma$ -behavioral equivalence of  $a, a' \in A_s$  by  $a \equiv_\Sigma^\Gamma a'$  iff  $A_c(a) = A_c(a')$  for all  $\Gamma$ -experiments  $c$ ; we may write  $\equiv$  instead of  $\equiv_\Sigma^\Gamma$  when  $\Sigma$  and  $\Gamma$  can be inferred from context. We require that all operations in  $\Sigma$  are compatible with  $\equiv_\Sigma^\Gamma$ . Note that behavioral equivalence is the identity on visible sorts, since the trivial contexts  $* : v$  are experiments for all  $v \in V$ . A major result in hidden logics, underlying the foundations of coinduction, is that  $\Gamma$ -behavioral equivalence is the largest equivalence which is identity on visible sorts and which is compatible with the operations in  $\Gamma$ .

Behavioral satisfaction of equations can now be naturally defined in terms of behavioral equivalence. A hidden  $\Sigma$ -algebra  $A$   $\Gamma$ -behaviorally satisfies a  $\Sigma$ -equation  $(\forall X) t = t'$ , say  $e$ , iff for each  $\theta : X \rightarrow A$ ,  $\theta(t) \equiv_\Sigma^\Gamma \theta(t')$ ; in this case we write  $A \models_\Sigma^\Gamma e$ . If  $E$  is a set of  $\Sigma$ -equations we then write  $A \models_\Sigma^\Gamma E$  when  $A$   $\Gamma$ -behaviorally satisfies each  $\Sigma$ -equation in  $E$ . We may omit  $\Sigma$  and/or  $\Gamma$  from  $\models_\Sigma^\Gamma$  when they are clear.

A behavioral  $\Sigma$ -specification is a triple  $(\Sigma, \Gamma, E)$  where  $\Sigma$  is a hidden signature,  $\Gamma$  is a hidden subsignature of  $\Sigma$ , and  $E$  is a set of  $\Sigma$ -sentences equations. Non-data  $\Gamma$ -operations (i.e., in  $\Gamma - \Sigma|_V$ ) are called *behavioral*. A  $\Sigma$ -algebra  $A$  *behaviorally satisfies* a behavioral specification  $\mathcal{B} = (\Sigma, \Gamma, E)$  iff  $A \models_\Sigma^\Gamma E$ , in which case we write  $A \models \mathcal{B}$ ; also  $\mathcal{B} \models e$  iff  $A \models \mathcal{B}$  implies  $A \models_\Sigma^\Gamma e$ .

EREs can be very naturally defined as a behavioral specification. The enormous benefit of doing so is that the behavioral inference, including most importantly coinduction, provide a *decision procedure* for equivalence of EREs. [46] shows how standard regular expressions (without negation) can be defined as a behavioral specification, a BOBJ implementation, and also how BOBJ with its circular coinductive rewriting algorithm can prove automatically several equivalences of regular expressions. Related interesting work can also be found in [110]. In this paper we extend that to general EREs, generate minimal observer monitors, and also give several other

examples.

**Example 7** *A behavioral specification of EREs defines a set of two visible sorts  $V = \{Bool, Event\}$ , one hidden sort  $H = \{Ere\}$ , one behavioral attribute  $\epsilon \in \_ : Ere \rightarrow Bool$  and one behavioral method, the derivative,  $\_[-] : Ere \times Event \rightarrow Ere$ , together with all the other operations in Subsection 7.2.2 defining EREs, including the events in  $E$  which are defined as visible constants of sort  $Event$ , and all the equations in Subsection 7.2.2. We call it the ERE behavioral specification and let  $\mathcal{B}_{ERE}$  denote it.*

*Since the only behavioral operators are the test for  $\epsilon$  membership and the derivative, it follows that the experiments have exactly the form  $\epsilon \in * \{a_1\} \{a_2\} \dots \{a_n\}$ , for any events  $a_1, a_2, \dots, a_n$ . In other words, an experiment consists of a series of derivations followed by an  $\epsilon$  membership test, and therefore two regular expressions are behavioral equivalent if and only if they cannot be distinguished by such experiments. Notice that the above reasoning applies within any algebra satisfying the presented behavioral specification. The one we are interested in is, of course, the free one, whose set carriers contain exactly the extended regular expressions as presented in Subsection 7.2.2, and the operations have the obvious interpretations. We informally call it the ERE algebra.*

Letting  $\equiv$  denote the behavioral equivalence relation generated on the ERE algebra, then Theorem 10 immediately yields the following important result.

**Theorem 11** *If  $R_1$  and  $R_2$  are two EREs then  $R_1 \equiv R_2$  if and only if  $\mathcal{L}(R_1) = \mathcal{L}(R_2)$ .*

This theorem allows us to prove equivalence of EREs by making use of behavioral inference in the ERE behavioral specification, from now on simply referred to by  $\mathcal{B}$ , including (especially) circular coinduction. The next section shows how circular coinduction works and how it can be used to show EREs equivalent.

#### *Circular Coinduction as an Inference Rule*

In the simplified version of hidden logics defined above, the usual equational inference rules, i.e., reflexivity, symmetry, transitivity, substitution and congruence [102] are all sound for behavioral satisfaction. However, equational reasoning can derive only a very limited amount of interesting behavioral equalities. For that reason, *circular coinduction* has been developed

as a very powerful automated technique to show behavioral equivalence. We let  $\Vdash$  denote the relation being defined by the equational rules plus circular coinduction, for deduction from a specification to an equation.

Before we present circular coinduction formally, we give the reader some intuitions by duality to structural induction. The reader who is only interested in using the presented procedure or who is not familiar with structural induction, can skip this paragraph. Inductive proofs show equality of terms  $t(x), t'(x)$  over a given variable  $x$  (seen as a constant) by showing  $t(\sigma(x))$  equals  $t'(\sigma(x))$  for all  $\sigma$  in a basis, while circular coinduction shows terms  $t, t'$  behaviorally equivalent by showing equivalence of  $\delta(t)$  and  $\delta(t')$  for all behavioral operations  $\delta$ . Coinduction applies behavioral operations at the top, while structural induction applies generator/constructor operations at the bottom. Both induction and circular coinduction assume some “frozen” instances of  $t, t'$  equal when checking the inductive/coinductive step: for induction, the terms are frozen at the bottom by replacing the induction variable by a constant, so that no other terms can be placed beneath the induction variable, while for coinduction, the terms are frozen at the top, so that they cannot be used as subterms of other terms (with some important but subtle exceptions which are not needed here; see [47]).

Freezing terms at the top is elegantly handled by a simple trick. Suppose every specification has a special visible sort  $b$ , and for each (hidden or visible) sort  $s$  in the specification, a special operation  $[-] : s \rightarrow b$ . No equations are assumed for these operations and no user defined sentence can refer to them; they are there for technical reasons. Thus, with just the equational inference rules, for any behavioral specification  $\mathcal{B}$  and any equation  $(\forall X) t = t'$ , it is necessarily the case that  $\mathcal{B} \Vdash (\forall X) t = t'$  iff  $\mathcal{B} \Vdash (\forall X) [t] = [t']$ . The rule below preserves this property. Let the sort of  $t, t'$  be hidden; then

**Circular Coinduction:**

$$\frac{\mathcal{B} \cup \{(\forall X) [t] = [t']\} \Vdash (\forall X, W) [\delta(t, W)] = [\delta(t', W)], \text{ for all appropriate } \delta \in \Gamma}{\mathcal{B} \Vdash (\forall X) t = t'}$$

We call the equation  $(\forall X) [t] = [t']$  added to  $\mathcal{B}$  a **circularity**; it could just as well have been called a coinduction hypothesis or a co-hypothesis, but we find the first name more intuitive because from a coalgebraic point of view, coinduction is all about finding circularities.

**Theorem 12** *The usual equational inference rules together with Circular Coinduction are sound. That means that if  $\mathcal{B} \Vdash (\forall X) t = t'$  and  $\text{sort}(t, t') \neq b$ , or if  $\mathcal{B} \Vdash (\forall X) [t] = [t']$ , then  $\mathcal{B} \models (\forall X) t = t'$ .*

**Example 8** *Suppose that we want to show that the EREs  $(a + b)^*$  and  $(a^*b^*)^*$  admit the same language. By Theorem 11, we can instead show that  $\mathcal{B}_{ERE} \models (\forall \emptyset) (a + b)^* = (a^*b^*)^*$ . Notice that  $a$  and  $b$  are treated as constant events here; one can also prove the result when  $a$  and  $b$  are variables, but one would need to first make use of the theorem of hidden constants [102]. To simplify writing, we omit the empty quantifier of equations. By the Circular Coinduction rule, one generates the following three proof obligations*

$$\begin{aligned} \mathcal{B}_{ERE} \cup \{[(a + b)^*] = [(a^*b^*)^*]\} &\Vdash [\epsilon \in (a + b)^*] = [\epsilon \in (a^*b^*)^*], \\ \mathcal{B}_{ERE} \cup \{[(a + b)^*] = [(a^*b^*)^*]\} &\Vdash [(a + b)^*\{a\}] = [(a^*b^*)^*\{a\}], \\ \mathcal{B}_{ERE} \cup \{[(a + b)^*] = [(a^*b^*)^*]\} &\Vdash [(a + b)^*\{b\}] = [(a^*b^*)^*\{b\}]. \end{aligned}$$

*The first proof task follows immediately by using the equations in  $\mathcal{B}$  as rewriting rules, while the other two tasks reduce to*

$$\begin{aligned} \mathcal{B}_{ERE} \cup \{[(a + b)^*] = [(a^*b^*)^*]\} &\Vdash [(a + b)^*] = [a^*(a^*b^*)^*], \\ \mathcal{B}_{ERE} \cup \{[(a + b)^*] = [(a^*b^*)^*]\} &\Vdash [(a + b)^*] = [b^*(a^*b^*)^*]. \end{aligned}$$

*By applying Circular Coinduction twice, after simplifying the two obvious proof tasks stating the  $\epsilon$  membership, one gets the following four proof obligations*

$$\begin{aligned} \mathcal{B}_{ERE} \cup \{[(a + b)^*] = [(a^*b^*)^*], [(a + b)^*] = [a^*(a^*b^*)^*]\} &\Vdash [(a + b)^*\{a\}] = [a^*(a^*b^*)^*\{a\}], \\ \mathcal{B}_{ERE} \cup \{[(a + b)^*] = [(a^*b^*)^*], [(a + b)^*] = [a^*(a^*b^*)^*]\} &\Vdash [(a + b)^*\{b\}] = [a^*(a^*b^*)^*\{b\}], \\ \mathcal{B}_{ERE} \cup \{[(a + b)^*] = [(a^*b^*)^*], [(a + b)^*] = [b^*(a^*b^*)^*]\} &\Vdash [(a + b)^*\{a\}] = [b^*(a^*b^*)^*\{a\}], \\ \mathcal{B}_{ERE} \cup \{[(a + b)^*] = [(a^*b^*)^*], [(a + b)^*] = [b^*(a^*b^*)^*]\} &\Vdash [(a + b)^*\{b\}] = [b^*(a^*b^*)^*\{b\}], \end{aligned}$$

*which, after simplification translate into*

$$\begin{aligned} \mathcal{B}_{ERE} \cup \{[(a + b)^*] = [(a^*b^*)^*], [(a + b)^*] = [a^*(a^*b^*)^*]\} &\Vdash [(a + b)^*] = [a^*(a^*b^*)^*], \\ \mathcal{B}_{ERE} \cup \{[(a + b)^*] = [(a^*b^*)^*], [(a + b)^*] = [a^*(a^*b^*)^*]\} &\Vdash [(a + b)^*] = [b^*(a^*b^*)^*], \\ \mathcal{B}_{ERE} \cup \{[(a + b)^*] = [(a^*b^*)^*], [(a + b)^*] = [b^*(a^*b^*)^*]\} &\Vdash [(a + b)^*] = [a^*(a^*b^*)^*], \\ \mathcal{B}_{ERE} \cup \{[(a + b)^*] = [(a^*b^*)^*], [(a + b)^*] = [b^*(a^*b^*)^*]\} &\Vdash [(a + b)^*] = [b^*(a^*b^*)^*], \end{aligned}$$

*Again by applying circular coinduction we get*

$$\begin{array}{l}
\mathcal{B}_{ERE} \cup \{[(a+b)^*] = [(a^*b^*)^*], [(a+b)^*] = [a^*(a^*b^*)^*], [(a+b)^*] = [b^*(a^*b^*)^*]\} \Vdash \\
\mathcal{B}_{ERE} \cup \{[(a+b)^*] = [(a^*b^*)^*], [(a+b)^*] = [a^*(a^*b^*)^*], [(a+b)^*] = [b^*(a^*b^*)^*]\} \Vdash \\
\mathcal{B}_{ERE} \cup \{[(a+b)^*] = [(a^*b^*)^*], [(a+b)^*] = [b^*(a^*b^*)^*], [(a+b)^*] = [a^*(a^*b^*)^*]\} \Vdash \\
\mathcal{B}_{ERE} \cup \{[(a+b)^*] = [(a^*b^*)^*], [(a+b)^*] = [b^*(a^*b^*)^*], [(a+b)^*] = [a^*(a^*b^*)^*]\} \Vdash
\end{array}$$

which now follow all immediately. Notice that *BOBJ* uses the newly added (to  $\mathcal{B}_{ERE}$ ) equations as rewriting rules when it applies its circular coinductive rewriting algorithm, so the proof above is done slightly differently, but entirely automatically.

**Example 9** Suppose now that one wants to show that  $\neg(a^*b) \equiv \epsilon + a^* + (a+b)^*b(a+b)(a+b)^*$ . One can also do it entirely automatically by circular coinduction as above, generating the following list of circularities:

$$\begin{array}{lcl}
[\neg(a^*b)] & = & [\epsilon + a^* + (a+b)^*b(a+b)(a+b)^*], \\
[\neg(\epsilon)] & = & [(a+b)^*b(a+b)(a+b)^* + (a+b)(a+b)^*], \\
[\neg(\emptyset)] & = & [(a+b)^*b(a+b)(a+b)^* + (a+b)^*], \\
[\neg(\emptyset)] & = & [(a+b)^*b(a+b)(a+b)^* + (a+b)(a+b)^* + (a+b)^*].
\end{array}$$

**Example 10** One can also show by circular coinduction that concrete EREs satisfy systems of guarded equations. This is an interesting but unrelated subject, so we do not discuss it in depth here. However, we show how easily one can prove by coinduction that  $a^*b$  is the solution of the equation  $R = a \cdot R + b$ . This equation can be given by adding a new ERE constant  $r$  to  $\mathcal{B}_{ERE}$ , together with the equations  $\epsilon \in r = \text{false}$ ,  $r\{a\} = r$ , and  $r\{b\} = \epsilon$ . Circular Coinduction applied on the goal  $r = a^*b$  generates the proof tasks:

$$\begin{array}{lcl}
\mathcal{B}_{ERE} \cup \{[r] = [a^*b]\} & \Vdash & [\epsilon \in r] = [\epsilon \in a^*b], \\
\mathcal{B}_{ERE} \cup \{[r] = [a^*b]\} & \Vdash & [r\{a\}] = [a^*b\{a\}], \\
\mathcal{B}_{ERE} \cup \{[r] = [a^*b]\} & \Vdash & [r\{b\}] = [a^*b\{b\}],
\end{array}$$

which all follow immediately.

The following says that circular coinduction provides a decision procedure for equivalence of EREs.

**Theorem 13** *If  $R_1$  and  $R_2$  are two EREs, then  $\mathcal{L}(R_1) = \mathcal{L}(R_2)$  if and only if  $\mathcal{B}_{ERE} \Vdash R_1 = R_2$ . Moreover, since the rules in  $\mathcal{B}_{ERE}$  are ground Church-Rosser and terminating, circular coinductive rewriting [46, 47], which iteratively rewrites proof tasks to their normal forms followed by a one step coinduction if needed, gives a decision procedure for ERE equivalence.*

### 7.2.5 Generating Minimal DFA Monitors by Coinduction

In this section we show how one can use the set of circularities generated by applying the circular coinduction rules in order to generate a minimal DFA from any ERE. This DFA can then be used as an optimal monitor for that ERE. The main idea here is to associate states in DFA to EREs obtained by deriving the initial ERE; when a new ERE is generated, it is tested for equivalence with all the other already generated EREs by using the coinductive procedure presented in the previous section. A crucial observation which significantly reduces the complexity of our procedure is that, once an equivalence is proved by circular coinductive rewriting, the entire set of circularities accumulated represent equivalent EREs. These can be used to later quickly infer the other equivalences, without having to generate the same circularities over and over again.

Since BOBJ does not (yet) provide any mechanism to return the set of circularities accumulated after proving a given behavioral equivalence, we were unable to use BOBJ to implement our optimal monitor generator. Instead, we have implemented our own version of coinductive rewriting engine for EREs, which is described below.

We are given an initial ERE  $R_0$  over alphabet  $A$  and from that we want to generate the equivalent minimal DFA  $D = (S, A, \delta, s_0, F)$ , where  $S$  is the set of states,  $\delta : S \times A \rightarrow S$  is the transition function,  $s_0$  is the initial state, and  $F \subseteq S$  is the set of final states. The coinductive rewriting engine explicitly accumulates the proven circularities in a set. The set is initialized to an empty set at the beginning of the algorithm. It is updated with the accumulated circularities whenever we prove equivalence of two regular expressions in the algorithm. The algorithm maintains the set of states  $S$  in the form of non-equivalent EREs. At the beginning of the algorithm  $S$  is initialized with a single element, which is the given ERE  $R_0$ . Next, we generate all the derivatives of the initial ERE one by one in a depth first manner. A derivative  $R_x = R\{x\}$  is added to the set  $S$ , if the set does not contain any ERE equivalent to the derivative  $R_x$ . We then extend the transition function by setting  $\delta(R, x) = R_x$ . If an ERE  $R'$  equivalent to the

derivative already exists in the set  $S$ , we extend the transition function by setting  $\delta(R, x) = R'$ . To check if an ERE equivalent to the derivative  $R_x$  already exists in the set  $S$ , we sequentially go through all the elements of the set  $S$  and try to prove its equivalence with  $R_x$ . In testing the equivalence we first add the set of circularities to the initial  $\mathcal{B}$ . Then we invoke the coinductive procedure. If for some ERE  $R' \in S$ , we are able to prove that  $R' \equiv R_x$  i.e  $\mathcal{B} \cup Eq_{\text{all}} \cup Eq_{\text{new}} \Vdash R' = R_x$ , then we add the new equivalences  $Eq_{\text{new}}$ , created by the coinductive procedure, to the set of circularities. Thus we reuse the already proven equivalences in future proofs.

The derivatives of the initial ERE  $R_0$  with respect to all events in the alphabet  $A$  are generated in a depth first fashion. The pseudo code for the whole algorithm is given in Figure 1.

```

dfs( $R$ )
begin
  foreach  $x \in A$  do
     $R_x \leftarrow R\{x\}$ ;
    if  $\exists R' \in S$  such that  $\mathcal{B} \cup Eq_{\text{all}} \cup Eq_{\text{new}} \Vdash R' = R_x$  then
       $\delta(R, x) = R'$ ;  $Eq_{\text{all}} \leftarrow Eq_{\text{all}} \cup Eq_{\text{new}}$ 
    else  $S \leftarrow S \cup \{R_x\}$ ;  $\delta(R, x) = R_x$ ; dfs( $R_x$ ); fi
  endfor
end

```

Figure 7.1: ERE to minimal DFA generation algorithm

In the procedure **dfs** the set of final states  $F$  consists of the EREs from  $S$  which contain  $\epsilon$ . This can be tested efficiently using the equations (10-16) in Subsection 7.2.2. The DFA generated by the procedure **dfs** may now contain some states which are non-final and from which the DFA can never reach a final state. We remove these redundant states by doing a breadth first search in backward direction from the final states. This can be done in time linear in the size of the DFA.

**Theorem 14** *If  $D$  is the DFA generated for a given ERE  $R$  by the above algorithm then*

1.  $\mathcal{L}(D) = \mathcal{L}(R)$ ,
2.  $D$  is the minimal DFA accepting  $\mathcal{L}(R)$ .

**Proof:** Suppose  $a_1a_2 \dots a_n \in \mathcal{L}(R)$ . Then  $\epsilon \in R\{a_1\}\{a_2\} \dots \{a_n\}$ . If  $R_i = R\{a_1\}\{a_2\} \dots \{a_i\}$  then  $R_{i+1} = R_i\{a_{i+1}\}$ . To prove that  $a_1a_2 \dots a_n \in \mathcal{L}(D)$ , we use induction to show that for each  $1 \leq i \leq n$ ,  $R_i \equiv \delta(R, a_1a_2 \dots a_i)$ . For the base case if  $R_1 \equiv R\{a_1\}$  then **dfs** extends the transition function by setting  $\delta(R, a_1) = R$ . Therefore,  $R_1 \equiv R = \delta(R, a_1)$ . If  $R_1 \not\equiv R$  then **dfs** extends  $\delta$  by setting  $\delta(R, a_1) = R_1$ . So  $R_1 \equiv \delta(R, a_1)$  holds in this case also. For the induction step let us assume that  $R_i \equiv R' = \delta(R, a_1a_2 \dots a_i)$ . If  $\delta(R', a_{i+1}) = R''$  then from the **dfs** procedure we can see that  $R'' \equiv R'\{a_{i+1}\}$ . However,  $R_i\{a_{i+1}\} \equiv R'\{a_{i+1}\}$ . So  $R_{i+1} \equiv R'' = \delta(R', a_{i+1}) = \delta(R, a_1a_2 \dots a_{i+1})$ . Also notice  $\epsilon \in R_n \equiv \delta(R, a_1a_2 \dots a_n)$ ; this implies that  $\delta(R, a_1a_2 \dots a_n)$  is a final state and hence  $a_1a_2 \dots a_n \in \mathcal{L}(D)$ .

Now suppose  $a_1a_2 \dots a_n \in \mathcal{L}(D)$ . The proof that  $a_1a_2 \dots a_n \in \mathcal{L}(R)$  goes in a similar way by showing that  $R_i \equiv \delta(R, a_1, a_2 \dots a_i)$ .  $\square$

## 7.2.6 Implementation and Evaluation

We have implemented the coinductive rewriting engine in the rewriting specification language Maude 2.0 [23]. The interested readers can download the implementation from the website <http://fsl.cs.uiuc.edu/rv/>. The operations on extended regular languages that are supported by our implementation are  $\sim$  for negation,  $*$  for Kleene Closure,  $_$  for concatenation,  $\&$  for intersection, and  $+$  for union in increasing order of precedence. Here, the intersection operator  $\&$  is a syntactic sugar and is translated to an ERE containing union and negation using De Morgan's Law:

$$\text{eq } R1 \ \& \ R2 = \sim (\sim R1 + \sim R2) \ .$$

To evaluate the performance of the algorithm we have generated the minimal DFA for all possible EREs of size up to 9. Surprisingly, the size of any DFA for EREs of size up to 9 did not exceed 9. Here the number of states gives the size of a DFA. The following table shows the performance of our procedure for the worst EREs of a given size. The code is executed on a

Pentium 4 2.4GHz, 4 GB RAM linux machine.

Size	ERE	no. of states in DFA	Time (ms)	Rewrites
4	$\neg (a \ b)$	4	< 1	863
5	$(a \neg b)^*$	4	< 1	1370
6	$\neg ((a \neg b)^*)$	4	1	1453
7	$\neg (a \neg a \ a)$	6	1	2261
8	$\neg ((a \neg b)^* \ b)$	7	1	3778
9	$\neg (a \neg a \ b) \ b$	9	5	9717

**Example 11** In particular, for the ERE  $\neg (a \neg a \ b) \ b$  the generated minimal DFA is given in Figure 7.2.

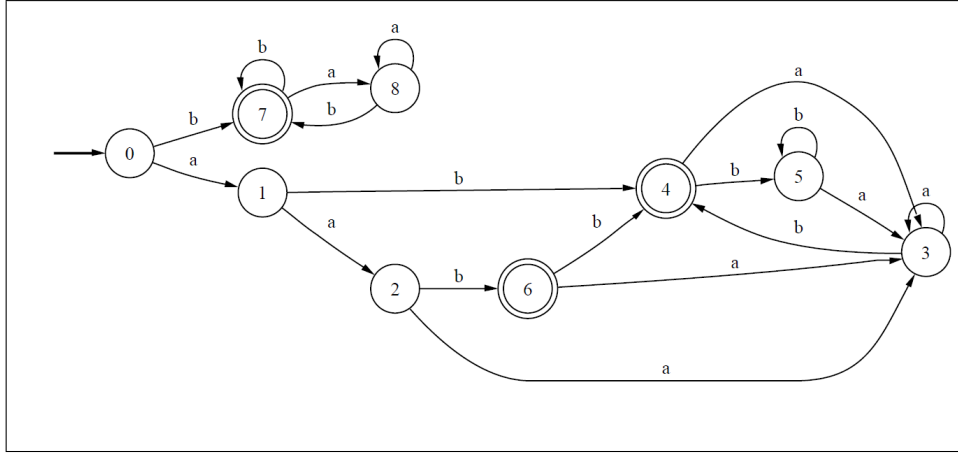


Figure 7.2:  $\neg (a \neg a \ b) \ b$

**Example 12** The ERE  $\neg ((\neg \text{empty}) (\text{green red}) (\neg \text{empty}))$  states the safety property that it should not be the case that in any trace of a traffic light we see green and red consecutively at any point. The set of events are assumed to be  $\{\text{green}, \text{red}, \text{yellow}\}$ . We think that this is the most intuitive and natural expression for this safety property. The implementation took 1ms and 1663 rewrites to generate the minimal DFA with 2 states. The DFA is given in Figure 7.3.

However for large EREs the algorithm may take a long time to generate a minimal DFA. The size of the generated DFA may grow non-elementarily

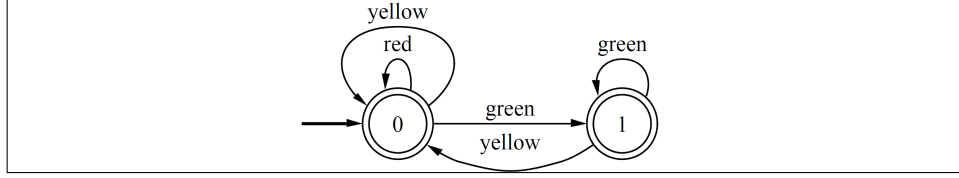


Figure 7.3:  $\neg ((\neg \text{empty}) (\text{green red}) (\neg \text{empty}))$

in the worst case. We generated DFAs for some complex EREs of larger sizes and got relatively promising results. One such sample result is as follows.

**Example 13** *Let us consider the following ERE of size 110*

$$\begin{aligned}
 &(\neg \$)^* \$ (\neg \$)^* \cap \\
 & (0 + 1 + \#)^* \# ( \\
 & \quad ((0 + 1)0\#(0 + 1 + \#)^* \$ (0 + 1)0 + (0 + 1)1\#(0 + 1 + \#)^* \$ (0 + 1)1) \\
 & \quad \cap (0(0 + 1)\#(0 + 1 + \#)^* \$ 0(0 + 1) + 1(0 + 1)\#(0 + 1 + \#)^* \$ 1(0 + 1))) .
 \end{aligned}$$

*This ERE accepts the language  $L_2$ , where*

$$L_k = \{\sigma \# w \# \sigma' \$ w \mid w \in \{0, 1\}^k \text{ and } \sigma, \sigma' \in \{0, 1, \#\}^*\}$$

*The language  $L_k$  was first introduced in [16] to show the power of alternation, used in [103] to show an exponential lower bound on ERE monitoring, and in [79, 80] to show the lower bounds for model checking. Our implementation took almost 18 minutes to generate the minimal DFA of size 107 and in the process it performed 1,374,089,220 rewrites.*

The above example shows that the procedure can take a large amount of time and space to generate DFAs for large EREs. To avoid the computation associated with the generation of minimal DFA we plan to maintain a database of EREs and their corresponding minimal DFAs on the internet. Whenever someone wants to generate the minimal DFA for a given ERE he/she can look up the internet database for the minimal DFA. If the ERE and the corresponding DFA exists in the database he/she can retrieve the corresponding DFA and use it as a monitor. Otherwise, he/she can generate the minimal DFA for the ERE and submit it to the internet database to create a new entry. The database will check the equivalence of the submitted ERE and the corresponding minimal DFA and insert it in the database. In this way one can avoid the computation of generating minimal DFA if it is already done by someone else. To further reduce the computation, circularities could also be stored in the database.

## Online Monitor Generation and Visualization

We have extended our implementation to create an internet server for optimal monitor generation that can be accessed from the url <http://fsl.cs.uiuc.edu/rv/>. Given an ERE the server generates the optimal DFA monitor for a user. The user submits the ERE through a web based form. A CGI script handling the web form takes the submitted ERE as an input, invokes the Maude implementation to generate the minimal DFA, and presents it to the user either as a graphical or a textual representation. To generate the graphical representation of the DFA we are currently using the GraphViz tool [39].

---

*We presented a new technique to generate optimal monitors for extended regular expressions, which avoids the traditional technique based on complementation of automata, that we think is quite complex and not necessary. Instead, we have considered the (co)algebraic definition of EREs and applied coinductive inferencing techniques in an innovative way to generate the minimal DFA. Our approach to store already proven equivalences has resulted into a very efficient and straightforward algorithm to generate minimal DFA. We have evaluated our implementation on several hundreds EREs and have got promising results in terms of running time. Finally we have installed a server on the internet which can generate the optimal DFA for a given ERE.*

*At least two major contributions have been made. Firstly, we have shown that coinduction is a viable and quite practical method to prove equivalence of extended regular expressions. Previously this was done only for regular expressions without complementation. Secondly, building on the coinductive technique, we have devised an algorithm to generate minimal DFAs from EREs. At present we have no bound for the size of the optimal DFA, but we know for sure that the DFAs we generate are indeed optimal. However we know that the size of an optimal DFA is bounded by some exponential in the size of the ERE. As future work, it seems interesting to investigate the size of minimal DFAs generated from EREs, and also to apply our coinductive techniques to generate monitors for other logics, such as temporal logics.*

---

## Chapter 8

# Monitor Synthesis: Linear Temporal Logic (LTL)

### 8.1 Finite Trace Future Time Linear Temporal Logic

Classical (infinite trace) linear temporal logic (LTL) [100, 88, 90] provides in addition to the propositional logic operators the temporal operators  $\Box$  (always),  $\Diamond$  (eventually),  $\Box U$  (until), and  $\circ$  (next). An LTL standard model is a function  $t : \mathcal{N}^+ \rightarrow 2^{\mathcal{P}}$  for some set of atomic propositions  $\mathcal{P}$ , i.e., an infinite trace over the alphabet  $2^{\mathcal{P}}$ , which maps each time point (a natural number) into the set of propositions that hold at that point. The operators have the following interpretation on such an infinite trace. Assume formulae  $X$  and  $Y$ . The formula  $\Box X$  holds if  $X$  holds in all time points, while  $\Diamond X$  holds if  $X$  holds in some future time point. The formula  $X U Y$  ( $X$  until  $Y$ ) holds if  $Y$  holds in some future time point, and until then  $X$  holds (so we consider strict until). Finally,  $\circ X$  holds for a trace if  $X$  holds in the suffix trace starting in the next (the second) time point. The propositional operators have their obvious meaning. For example, the formula  $\Box (X \rightarrow \Diamond Y)$  is true if for any time point ( $\Box$ ) it holds that if  $X$  is true then eventually ( $\Diamond$ )  $Y$  is true. It is standard to define a core LTL using only atomic propositions, the propositional operators  $\neg$  (not) and  $\wedge$  (and), and the temporal operators  $\circ$  and  $\Box U$ , and then define all other propositional and temporal operators as derived constructs. Standard equations are  $\Diamond X = \text{true} U X$  and  $\Box X = \neg \Diamond \neg X$ .

Our goal is to develop a framework for testing software systems using temporal logic. Tests are performed on finite execution traces and we

therefore need to formalize what it means for a *finite trace* to satisfy an LTL formula. We first present a semantics of finite trace LTL using standard mathematical notation. Then we present a specification in Maude of a finite trace semantics. Whereas the former semantics uses universal and existential quantification, the second Maude specification is defined using recursive definitions that have a straightforward operational rewriting interpretation and which therefore can be executed.

### 8.1.1 Finite Trace Semantics

As mentioned in Subsection 2.1.2, a trace is viewed as a non-empty finite sequence of program states, each state denoting the set of propositions that hold at that state. We shall first outline the finite trace LTL semantics using standard mathematical notation rather than Maude notation. The debatable issue here is what happens at the end of the trace. The choice to validate or invalidate all the atomic propositions does not work in practice, because there might be propositions whose values are always opposite to each other, such as, for example, “gate up” and “gate down”. Driven by experiments, we found that a more reasonable assumption is to regard a finite trace as an infinite stationary trace in which the last event is repeated infinitely.

Assume two total functions on traces,  $head : \text{Trace} \rightarrow \text{Event}$  returning the head event of a trace and  $length$  returning the length of a finite trace, and a partial function  $tail : \text{Trace} \rightarrow \text{Trace}$  for taking the tail of a trace. That is,  $head(e, t) = head(e) = e$ ,  $tail(e, t) = t$ , and  $length(e) = 1$  and  $length(e, t) = 1 + length(t)$ . Assume further for any trace  $t$ , that  $t_i$  denotes the suffix trace that starts at position  $i$ , with positions starting at 1. The satisfaction relation  $\models \subseteq \text{Trace} \times \text{Formula}$  defines when a trace  $t$  satisfies a formula  $f$ , written  $t \models f$ , and is defined inductively over the structure of the formulae as follows, where  $A$  is any atomic proposition and  $X$  and  $Y$  are any formulae:

$t \models \text{true}$	iff	$true$ ,
$t \models \text{false}$	iff	$false$ ,
$t \models A$	iff	$A \in head(t)$ ,
$t \models X \wedge Y$	iff	$t \models X$ and $t \models Y$ ,
$t \models X \vee Y$	iff	$t \models X$ or $t \models Y$ ,
$t \models \circ X$	iff	(if $tail(t)$ is defined then $tail(t) \models X$ else $t \models X$ ),
$t \models \langle \rangle X$	iff	$(\exists i \leq length(t)) t_i \models X$ ,
$t \models [] X$	iff	$(\forall i \leq length(t)) t_i \models X$ ,
$t \models X \cup Y$	iff	$(\exists i \leq length(t)) (t_i \models Y \text{ and } (\forall j < i) t_j \models X)$ .

The semantics of the “next” operator reflects perhaps best the stationarity assumption of last events in finite traces.

Notice that finite trace LTL can behave quite differently from standard infinite trace LTL. For example, there are formulae which are not valid in infinite trace LTL but valid in finite trace LTL, such as  $\langle \rangle (\Box A \vee \Box !A)$  for any atomic proposition  $A$ , and there are formulae which are satisfiable in infinite trace LTL and not satisfiable in finite trace LTL, such as the negation of the above. The formula above is satisfied by any finite trace because the last event/state in the trace either contains  $A$  or it does not.

### 8.1.2 Finite Trace Semantics in Maude

Now it can be relatively easily seen that the following Maude specification correctly “defines” the finite trace semantics of LTL described above. The only important deviation from the rigorous mathematical formulation described above is that the quantifiers over finite sets of indexes are expressed recursively.

```
fmod LTL is
  extending PROP-CALC .
*** syntax
  op []_ : Formula -> Formula [prec 11] .
  op <>_ : Formula -> Formula [prec 11] .
  op _U_ : Formula Formula -> Formula [prec 14] .
  op o_ : Formula -> Formula [prec 11] .
*** semantics
  vars X Y : Formula .
  var E : Event .
  var T : Trace .
  eq E |= o X = E |= X .
  eq E,T |= o X = T |= X .
  eq E |= <> X = E |= X .
  eq E,T |= <> X = E,T |= X or T |= <> X .
  eq E |= [] X = E |= X .
  eq E,T |= [] X = E,T |= X and T |= [] X .
  eq E |= X U Y = E |= Y .
  eq E,T |= X U Y = E,T |= Y or E,T |= X and T |= X U Y .
endfm
```

Notice that only the temporal operators needed declarations and semantics, the others being already defined in PROP-CALC and LOGICS-BASIC, and that the definitions that involved the functions *head* and *tail* were replaced by two alternative equations.

One can now directly verify LTL properties on finite traces using Maude's rewriting engine. Consider as an example a traffic light that switches between the colors *green*, *yellow*, and *red*. The LTL property that after *green* comes *yellow*, and its negation, can now be verified on a finite trace using Maude's rewriting engine, by typing commands to Maude such as:

```
reduce green, yellow, red, green, yellow, red, green, yellow, red, red
  |= [] (green -> !red U yellow) .
reduce green, yellow, red, green, yellow, red, green, yellow, red, red
  |= !([] (green -> !red U yellow)) .
```

which should return the expected answers, i.e., **true** and **false**, respectively. The algorithm above does nothing but blindly follows the mathematical definition of satisfaction and even runs reasonably fast for relatively small traces. For example, it takes<sup>1</sup> about 30ms (74k rewrite steps) to reduce the first formula above and less than 1s (254k rewrite steps) to reduce the second on traces of 100 events (10 times larger than the above). Unfortunately, this algorithm does not seem to be tractable for large event traces, even if run on very fast platforms. As a concrete practical example, it took Maude 7.3 million rewriting steps (3 seconds) to reduce the first formula above and 2.4 billion steps (1000 seconds) for the second on traces of 1,000 events; it could not finish in one night (more than 10 hours) the reduction of the second formula on a trace of 10,000 events. Since the event traces generated by an executing program can easily be larger than 10,000 events, the trivial algorithm above cannot be used in practice.

A rigorous complexity analysis of the algorithm above is hard (because it has to take into consideration the evaluation strategy used by Maude for terms of sort `Bool`) and not worth the effort. However, a simplified worse-case analysis can be easily made if one only counts the maximum number of atoms of the form `event |= atom` that can occur during the rewriting of a satisfaction term, as if all the Boolean reductions were applied after all the other reductions: let us consider a formula  $X = [] ([] (\dots ([] A) \dots))$  where the always operator is nested  $m$  times, and a trace  $T$  of size  $n$ , and let  $T(n, m)$  be the total number of basic satisfactions `event |= atom` that occur in the normal form of the term  $T \models X$  if no Boolean reductions were applied. Then, the recurrence formula  $T(n, m) = T(n - 1, m) + T(n, m - 1)$  follows immediately from the specification above. Since  $\binom{n}{m} = \binom{n-1}{m} + \binom{n-1}{m-1}$ , it follows that  $T(n, m) > \binom{n}{m}$ , that is,  $T(n, m) = \Omega(n^m)$ , which is of course unacceptable.

---

<sup>1</sup>On a 1.7GHz, 1Gb memory PC.

## 8.2 A Backwards, Asynchronous, but Efficient Algorithm

The satisfaction relation above for finite trace LTL can hence be defined recursively, both on the structure of the formulae and on the size of the execution trace. As is often the case for functions defined this way, an efficient *dynamic programming* algorithm can be generated from any LTL formula. We first show how such an algorithm looks for a particular formula, and then present the main algorithm generator. The work in this section appeared as a technical report [104], but for a slightly different finite trace LTL, namely one in which all the atomic propositions were considered *false* at the end of the trace. As explained previously in the paper, we are now in the favor of a semantics where traces are considered stationary in their last event. The generated dynamic programming algorithms are as efficient as they can be and one can hope: linear in both the trace and the LTL formula. Unfortunately, they need to traverse the execution trace backwards, so they are trace storing and asynchronous. However, a similar but dual technique applies to past time LTL, producing very efficient forwards and synchronous algorithms [60, 59].

### 8.2.1 An Example

The formula we choose below is artificial (and will not be used later in the paper), but contains all four temporal operators. We believe that this example would practically be sufficient for the reader to foresee the general algorithm presented in the remaining of the section. Let  $\Box((p \cup q) \rightarrow \langle \rangle (q \rightarrow or))$  be an LTL formula and let  $\varphi_1, \varphi_2, \dots, \varphi_{10}$  be its subformulae, in breadth-first order:

$$\begin{aligned}
 \varphi_1 &= \Box((p \cup q) \rightarrow \langle \rangle (q \rightarrow or)), \\
 \varphi_2 &= (p \cup q) \rightarrow \langle \rangle (q \rightarrow or), \\
 \varphi_3 &= p \cup q, \\
 \varphi_4 &= \langle \rangle (q \rightarrow or), \\
 \varphi_5 &= p, \\
 \varphi_6 &= q, \\
 \varphi_7 &= q \rightarrow or, \\
 \varphi_8 &= q, \\
 \varphi_9 &= or, \\
 \varphi_{10} &= r.
 \end{aligned}$$

Given any finite trace  $t = e_1 e_2 \dots e_n$  of  $n$  events, one can recursively define a matrix  $s[1..n, 1..10]$  of Boolean values  $\{0, 1\}$ , with the meaning that  $s[i, j] = 1$  iff  $t_i \models \varphi_j$  as follows:

$$\begin{aligned}
s[i, 10] &= (r \in e_i) \\
s[i, 9] &= s[i + 1, 10] \\
s[i, 8] &= (q \in e_i) \\
s[i, 7] &= s[i, 8] \text{ implies } s[i, 9] \\
s[i, 6] &= (q \in e_i) \\
s[i, 5] &= (p \in e_i) \\
s[i, 4] &= s[i, 7] \text{ or } s[i + 1, 4] \\
s[i, 3] &= s[i, 6] \text{ or } (s[i, 5] \text{ and } s[i + 1, 3]) \\
s[i, 2] &= s[i, 3] \text{ implies } s[i, 4] \\
s[i, 1] &= s[i, 2] \text{ and } s[i + 1, 1],
\end{aligned}$$

for all  $i < n$ , where *and*, *or*, *implies* are ordinary Boolean operations and  $==$  is the equality predicate, where  $s[n, 1..10]$  are defined as below:

$$\begin{aligned}
s[n, 10] &= (r \in e_n) \\
s[n, 9] &= s[n, 10] \\
s[n, 8] &= (q \in e_n) \\
s[n, 7] &= s[n, 8] \text{ implies } s[n, 9] \\
s[n, 6] &= (q \in e_n) \\
s[n, 5] &= (p \in e_n) \\
s[n, 4] &= s[n, 7] \\
s[n, 3] &= s[n, 6] \\
s[n, 2] &= s[n, 3] \text{ implies } s[n, 4] \\
s[n, 1] &= s[n, 2].
\end{aligned}$$

Note again that the trace needs to be *traversed backwards*, and that the row  $n$  of  $s$  is filled according to the stationary view of finite traces in their last event. An important observation is that, like in many other dynamic programming algorithms, one does not have to store all the table  $s[1..n, 1..10]$ , which would be quite large in practice; in this case, one needs only two rows,  $s[i, 1..10]$  and  $s[i + 1, 1..10]$ , which we shall write *now* and *next* from now on, respectively. It is now only a simple exercise to write up the following algorithm:

INPUT: trace  $t = e_1 e_2 \dots e_n$   
 $next[10] \leftarrow (r \in e_n);$

```

next[9]  $\leftarrow$  next[10];
next[8]  $\leftarrow$  ( $q \in e_n$ );
next[7]  $\leftarrow$  next[8] implies next[9];
next[6]  $\leftarrow$  ( $q \in e_n$ );
next[5]  $\leftarrow$  ( $p \in e_n$ );
next[4]  $\leftarrow$  next[7];
next[3]  $\leftarrow$  next[6];
next[2]  $\leftarrow$  next[3] implies next[4];
next[1]  $\leftarrow$  next[2];
for  $i = n - 1$  downto 1 do {
    now[10]  $\leftarrow$  ( $r \in e_i$ );
    now[9]  $\leftarrow$  next[10];
    now[8]  $\leftarrow$  ( $q \in e_i$ );
    now[7]  $\leftarrow$  now[8] implies now[9];
    now[6]  $\leftarrow$  ( $q \in e_i$ );
    now[5]  $\leftarrow$  ( $p \in e_i$ );
    now[4]  $\leftarrow$  now[7] or next[4];
    now[3]  $\leftarrow$  now[6] or (now[5] and next[3]);
    now[2]  $\leftarrow$  now[3] implies now[4];
    now[1]  $\leftarrow$  now[2] and next[1];
    next  $\leftarrow$  now }
output(next[1]);

```

The algorithm above can be further optimized, noticing that only the bits 10, 4, 3 and 1 are needed in the vectors *now* and *next*, as we did for past time LTL in [60, 59]. The analysis of this algorithm is straightforward. Its time complexity is  $\Theta(n \cdot m)$  while the memory required is  $2 \cdot m$  bits, where  $n$  is the length of the trace and  $m$  is the size of the LTL formula.

### 8.2.2 Generating Dynamic Programming Algorithms

We now formally describe our algorithm that synthesizes dynamic programming algorithms like the one above from LTL formulae. Our synthesizer is generic, the potential user being expected to adapt it to his/her desired target language. The algorithm consists of three main steps:

**Breadth First Search.** The LTL formula should be first visited in breadth-first search (BFS) order to assign increasing numbers to subformulae as they are visited. Let  $\varphi_1, \varphi_2, \dots, \varphi_m$  be the list of all subformulae in BFS order. Because of the semantics of finite trace LTL, this step ensures us that the truth value of  $t_i \models \varphi_j$  can be completely determined from the truth values of  $t_i \models \varphi_{j'}$  for all  $j < j' \leq m$  and the truth values of  $t_{i+1} \models \varphi_{j'}$  for all  $j \leq j' \leq m$ . This recurrence gives the order in which one should generate the code.

**Loop Initialization.** Before we generate the “for” loop, we should first initialize the vector  $next[1..m]$ , which basically gives the truth values of the subformulae on the empty trace. According to the semantics of LTL, one should fill the vector  $next$  backwards. For a given  $m \geq j \geq 1$ ,  $next[j]$  is calculated as follows:

- If  $\varphi_j$  is a variable then  $next[j] = (\varphi_j \in e_n)$ . In a more complex setting, where  $\varphi_j$  was a state predicate, one would have to evaluate  $\varphi_j$  in the final state in the execution trace;
- If  $\varphi_j$  is  $!\varphi_{j'}$  for some  $j < j' \leq m$ , then  $next[j] = not\ next[j']$ , where  $not$  is the negation operation on Booleans (bits);
- If  $\varphi_j$  is  $\varphi_{j_1} Op\ \varphi_{j_2}$  for some  $j < j_1, j_2 \leq m$ , then  $next[j] = next[j_1] op\ next[j_2]$ , where  $Op$  is any propositional operation and  $op$  is its corresponding Boolean operation;
- If  $\varphi_j$  is  $\circ\varphi_{j'}$ ,  $\square\varphi_{j'}$ , or  $<>\varphi_{j'}$ , then clearly  $next[j] = next[j']$  according to the stationary semantics of our finite trace LTL;
- If  $\varphi_j$  is  $\varphi_{j_1} \cup \varphi_{j_2}$  for some  $j < j_1, j_2 \leq m$ , then  $next[j] = next[j_2]$  for the same reason as above.

**Loop Generation.** Because of the dependences in the recursive definition of finite trace LTL satisfaction relation, one is expected to visit the remaining of the trace backwards, so the loop index will vary from  $n - 1$  down to 1. The loop body will update/calculate the vector  $now$  and in the end will move it into the vector  $next$  to serve as basis for the next iteration. At a certain iteration  $i$ , the vector  $now$  is updated also backwards as follows:

- If  $\varphi_j$  is a variable then  $now[j] = (\varphi_j \in e_i)$ .
- If  $\varphi_j$  is  $!\varphi_{j'}$  for some  $j < j' \leq m$ , then  $now[j] = not\ now[j']$ ;

- If  $\varphi_j$  is  $\varphi_{j_1} Op \varphi_{j_2}$  for  $j < j_1, j_2 \leq m$ , then  $now[j] = now[j_1] op now[j_2]$ , where  $Op$  is any propositional operation and  $op$  is its corresponding Boolean operation;
- If  $\varphi_j$  is  $\circ\varphi_{j'}$  then  $now[j] = next[j']$  since  $\varphi_j$  holds now if and only if  $\varphi_{j'}$  held at the previous step (which processed the next event, the  $i + 1$ -th);
- If  $\varphi_j$  is  $\Box\varphi_{j'}$  then  $now[j] = now[j']$  and  $next[j]$  because  $\varphi_j$  holds now if and only if  $\varphi_{j'}$  holds now and  $\varphi_j$  held at the previous iteration;
- If  $\varphi_j$  is  $\langle\rangle\varphi_{j'}$  then  $now[j] = now[j']$  or  $next[j]$  because of similar reasons as above;
- If  $\varphi_j$  is  $\varphi_{j_1} \cup \varphi_{j_2}$  for some  $j < j_1, j_2 \leq m$ , then  $now[j] = now[j_2]$  or  $(now[j_1] \text{ and } next[j])$ .

After each iteration,  $next[1]$  says whether the initial LTL formula is validated by the trace  $e_i e_{i+1} \dots e_n$ . Therefore, the desired output is  $next[1]$  after the last iteration. Putting all the above together, one can now write up the generic pseudocode presented below which can be implemented very efficiently on any current platform. Since the BFS procedure is linear, the algorithm synthesizes a dynamic programming algorithm from an LTL formula in linear time and of linear size with the size of the formula.

The following generic program implements the discussed technique. It takes as input an LTL formula and generates a “for” loop which traverses the trace of events backwards, thus validating or invalidating the formula.

INPUT: LTL formula  $\varphi$

**output**(“INPUT: trace  $t = e_1 e_2 \dots e_n$ ”);

let  $\varphi_1, \varphi_2, \dots, \varphi_m$  be all the subformulae of  $\varphi$  in BFS order

**for**  $j = m$  **downto** 1 **do** {

**output**(“ $next[$ ”,  $j$ , “]  $\leftarrow$  ”);

**if**  $\varphi_j$  is a variable **then** **output**(“( $\varphi_j$ , “ $\in e_n$ );”);

**if**  $\varphi_j = \neg\varphi_{j'}$  **then** **output**(“not  $next[$ ”,  $j'$ , “];”);

**if**  $\varphi_j = \varphi_{j_1} Op \varphi_{j_2}$  **then** **output**(“ $next[$ ”,  $j_1$ , “]  $op$   $next[$ ”,  $j_2$ , “];”);

**if**  $\varphi_j = \circ\varphi_{j'}$  **then** **output**(“ $next[$ ”,  $j'$ , “];”);

**if**  $\varphi_j = \Box\varphi_{j'}$  **then** **output**(“ $next[$ ”,  $j'$ , “];”);

**if**  $\varphi_j = \langle\rangle\varphi_{j'}$  **then** **output**(“ $next[$ ”,  $j'$ , “];”);

**if**  $\varphi_j = \varphi_{j_1} \cup \varphi_{j_2}$  **then** **output**(“ $next[$ ”,  $j_2$ , “];”); }

```

output("for  $i = n - 1$  downto 1 do {");
for  $j = m$  downto 1 do {
    output("     $now[j, " \leftarrow "$ ");
    if  $\varphi_j$  is a variable then output("( $\varphi_j, \in e_i$ );");
    if  $\varphi_j = !\varphi_{j'}$  then output("not  $now[j', "$ ");
    if  $\varphi_j = \varphi_{j_1} Op \varphi_{j_2}$  then output("now[j1, " op now[j2, "];");
    if  $\varphi_j = \circ\varphi_{j'}$  then output("next[j', "];");
    if  $\varphi_j = \square\varphi_{j'}$  then output("now[j', " and next[j, "];");
    if  $\varphi_j = \langle\rangle\varphi_{j'}$  then output("now[j', " or next[j, "];");
    if  $\varphi_j = \varphi_{j_1} \cup \varphi_{j_2}$  then output("now[j2, " or (now[j1, "
and next[j, "]);"); }
output("    next  $\leftarrow$  now; }");
output("output next[1];");

```

where  $Op$  is any propositional connective and  $op$  is its corresponding Boolean operator.

The Boolean operations used above are usually very efficiently implemented on any microprocessor and the vectors of bits  $next$  and  $now$  are small enough to be kept in cache. Moreover, the dependencies between instructions in the generated “for” loop are simple to analyze, so a reasonable compiler can easily unfold or/and parallelize it to take advantage of machine’s resources. Consequently, the generated code is expected to run very fast.

The dynamic programming technique presented in this section is as efficient as one can hope, but, unfortunately, has a major drawback: it needs to traverse the execution trace backwards. From a practical perspective, that means that the instrumented program is run for some period of time while its execution trace is saved, and then, after the program was stopped, its execution trace is traversed backwards and (efficiently) analyzed. Besides the obvious inconvenience due to storing potentially huge execution traces, this method cannot be used to monitor programs synchronously.

### 8.3 A Forwards and Often Synchronous Algorithm

In this section we shall present a more efficient rewriting semantics for LTL, based on the idea of consuming the events in the trace, one by one, and updating a data structure (which is also a formula) corresponding to the effect of the event on the value of the formula. An important advantage of

this algorithm is that it often detects when a formula is violated or validated before the end of the execution trace, so, unlike the algorithms above, it is suitable for online monitoring. Our decision to write an operational semantics this way was motivated by an attempt to program such an algorithm in Java, where such a solution would be natural. The presented rewriting-based algorithm is linear in the size of the execution trace and worst-case exponential in the size of the monitored LTL formula.

### 8.3.1 An Event Consuming Algorithm

We will implement this rewriting based algorithm by extending the definition of the event consuming operation  $\_ \{ \_ \} : \text{Formula Event}^* \rightarrow \text{Formula}$  to temporal operators, with the following intuition. Assuming a trace  $E, T$  consisting of event  $E$  followed by trace  $T$ , a formula  $X$  holds on this trace if and only if  $X\{E\}$  holds on the remaining trace  $T$ . If the event  $E$  is terminal then  $X\{E \ * \}$  holds if and only if  $X$  holds under standard LTL semantics on the infinite trace containing only the event  $E$ .

```
fmod LTL-REVISED is
  protecting LTL .
  vars X Y : Formula .
  var E : Event .
  var T : Trace .
  eq (o X){E} = X .
  eq (o X){E *} = X{E *} .
  eq (<> X){E} = X{E} /\ <> X .
  eq (<> X){E *} = X{E *} .
  eq ([] X){E} = X{E} /\ [] X .
  eq ([] X){E *} = X{E *} .
  eq (X U Y){E} = Y{E} /\ X{E} /\ X U Y .
  eq (X U Y){E *} = Y{E *} .

  op _|-_ : Trace Formula -> Bool .
  eq E |- X = [X{E *}] .
  eq E, T |- X = T |- X{E} .
endfm
```

The rule for the temporal operator  $[]X$  should be read as follows: the formula  $X$  must hold now ( $X\{E\}$ ) and also in the future ( $[]X$ ). The sub-expression  $X\{E\}$  represents the formula that must hold on the rest of the trace in order for  $X$  to hold now.

As an example, consider again the traffic light controller safety formula  $\Box(\text{green} \rightarrow !\text{red} \text{ U } \text{yellow})$ , which is first rewritten to  $\Box(\text{true} ++ \text{green} ++ \text{green} \wedge (\text{true} ++ \text{red}) \text{ U } \text{yellow})$  by the equations in module PROP-CALC. This formula modified by an event `green yellow` (notice that two lights can be lit at the same time) yields the rewriting sequence

```

( $\Box(\text{true} ++ \text{green} ++ \text{green} \wedge (\text{true} ++ \text{red}) \text{ U } \text{yellow})$ ){green yellow} ==>
(true ++ green{green yellow}
  ++ green{green yellow} /\ ((true ++ red) U yellow){green yellow}
  /\  $\Box(\text{true} ++ \text{green} ++ \text{green} \wedge (\text{true} ++ \text{red}) \text{ U } \text{yellow})$  ==>
((true ++ red) U yellow){green yellow}
  /\  $\Box(\text{true} ++ \text{green} ++ \text{green} \wedge (\text{true} ++ \text{red}) \text{ U } \text{yellow})$  ==>
(yellow{green yellow} \/ ((true ++ red{green yellow}) /\ (true ++ red) U yellow)
  /\  $\Box(\text{true} ++ \text{green} ++ \text{green} \wedge (\text{true} ++ \text{red}) \text{ U } \text{yellow})$  ==>
 $\Box(\text{true} ++ \text{green} ++ \text{green} \wedge (\text{true} ++ \text{red}) \text{ U } \text{yellow})$ 

```

which is exactly the original formula, while the same formula transformed by just the event `green` yields

```

( $\Box(\text{true} ++ \text{green} ++ \text{green} \wedge (\text{true} ++ \text{red}) \text{ U } \text{yellow})$ ){green} ==>
(true ++ green{green}
  ++ green{green} /\ ((true ++ red) U yellow){green}
  /\  $\Box(\text{true} ++ \text{green} ++ \text{green} \wedge (\text{true} ++ \text{red}) \text{ U } \text{yellow})$  ==>
((true ++ red) U yellow){green}
  /\  $\Box(\text{true} ++ \text{green} ++ \text{green} \wedge (\text{true} ++ \text{red}) \text{ U } \text{yellow})$  ==>
(yellow{green} \/ ((true ++ red{green}) /\ (true ++ red) U yellow)
  /\  $\Box(\text{true} ++ \text{green} ++ \text{green} \wedge (\text{true} ++ \text{red}) \text{ U } \text{yellow})$  ==>
(true ++ red) U yellow /\  $\Box(\text{true} ++ \text{green} ++ \text{green} \wedge (\text{true} ++ \text{red}) \text{ U } \text{yellow})$ 

```

which further modified by an event `red` yields

```

(yellow{red} \/ (true ++ red{red}) /\ (true ++ red) U yellow)
  /\ ( $\Box(\text{true} ++ \text{green} ++ \text{green} \wedge (\text{true} ++ \text{red}) \text{ U } \text{yellow})$ ){red} ==>
false /\ ( $\Box(\text{true} ++ \text{green} ++ \text{green} \wedge (\text{true} ++ \text{red}) \text{ U } \text{yellow})$ ){red} ==>
false

```

When the current formula becomes `false`, as it happened above, we say that the original formula has been violated. Indeed, the current formula will remain `false` for any subsequent trace of events, so the result of the monitoring session will be `false`.

Note that the rewriting system described so far obviously terminates, because what it does is to propagate the current event to the atomic subformulae, replace those by either `true` or `false`, and eventually canonize the newly obtained formula.

Some operators could be defined in terms of others, as is typically the case in the standard semantics for LTL. For example, we could introduce an equation of the form:  $\langle \rangle X = \text{true} \cup X$ , and then eliminate the rewriting rule for  $\langle \rangle X$  in the above module. This turns out to be less efficient in practice though, because more rewrites are needed. This happens regardless of whether one enables memoization (explained in detail in Subsection 8.3.3) or not, because memoization brings a real benefit only when previously processed terms are to be reduced again.

This module eventually defines a new satisfaction relation  $\_ \models \_$  between traces and formulae. The term  $T \models X$  is evaluated now by an iterative traversal over the trace, where each event transforms the formula. Note that the new formula that is generated at each step is always kept small by being reduced to normal form via the equations in the PROP-CALC module in Subsection 2.1.2.

Our current JPAX implementation of the rewriting algorithm above executes the last two rules of the module LTL-REVISED outside the main rewriting engine. More precisely, Maude is started in its *loop mode* [26], which provides the capability of enabling rewriting rules in a reactive system style: a “state” term is stored, which can then be modified via rewriting rules that are activated by ASCII text events that are provided via the standard I/O. JPAX starts a Maude process and assigns the formula to be monitored as its loop mode state. Then, as events are received from the monitored program, they are filtered and forwarded to the Maude module, which then enables rewriting on the term  $X\{E\}$ , where  $X$  is the current formula and  $E$  is the newly received event; the normal form of this reduction, a formula, is stored as the new loop mode state term. The process continues until the last event is received. JPAX tags the last event, asking Maude to reduce a term  $X\{E \ast\}$ ; the result will be either `true` or `false`, which is reported to the user at the end of the monitoring session<sup>2</sup>.

A natural question here is how big the stored formula can grow during a monitoring session. Such a formula will consist of Boolean combinations of sub-formulae of the initial formula, kept in a minimal canonical form. This can grow exponentially in the size of the initial formula in the worst-case (see [103] for a related result for extended regular expressions).

**Theorem 15** *For any formula  $X$  of size  $m$  and any sequence of events to be monitored  $E_1, E_2, \dots, E_n$ , the formula  $X\{E_1\}\{E_2\} \dots \{E_n\}$  needs  $O(2^m)$*

---

<sup>2</sup>In fact, JPAX reports a similar message also when the current monitoring requirement becomes `true` or `false` at any time during the monitoring process.

space to be stored. Moreover, the exponential space cannot be avoided: any synchronous or asynchronous forwards monitoring algorithm for LTL requires space  $\Omega(2^{c\sqrt{m}})$  space, where  $c$  is some fixed constant.

**Proof:** Due to the Boolean ring simplification rules in PROP-CALC, any LTL formula is kept in a canonical form, which is an exclusive disjunction of conjunctions, where conjuncts have temporal operators at top. Moreover, after processing any number of events, in our case  $E1, E2, \dots, En$ , the conjuncts in the normal form of  $X\{E1\}\{E2\} \dots \{En\}$  are subterms of the initial formula  $X$ , each having a temporal operator at its top. Since there are at most  $m$  such subformulae of  $X$ , it follows that there are at most  $2^m$  possibilities to combine them in a conjunction. Therefore, one needs space  $O(2^m)$  to store any exclusive disjunction of such conjunctions. This reasoning only applies on “idealistic” rewriting engines, which carefully optimize space needs during rewriting. It is not clear to us whether Maude is able to attain this space upper bound in all situations.

For the space lower bound of any finite trace LTL monitoring algorithm, consider a simplified framework with only two atomic predicate and therefore only four possible states. For simplicity, we encode these four states by 0, 1, # and \$. Consider also some natural number  $k$  and the language

$$L_k = \{\sigma \# w \# \sigma' \$ w \mid w \in \{0, 1\}^k \text{ and } \sigma, \sigma' \in \{0, 1, \#, \$\}^*\}.$$

This language was previously used in several works [79, 80, 103] to prove lower bounds. The language can be shown to contain exactly those finite traces satisfying the following LTL formula [80] of size  $\Theta(k^2)$ :

$$\phi_k = [(!\$)\mathcal{U}(\$ \wedge \circ \square(!\$))] \wedge \diamond[\# \wedge \circ^{n+1} \# \wedge \bigwedge_{i=1}^n ((\circ^i 0 \wedge \square(\$ \rightarrow \circ^i 0)) \vee (\circ^i 1 \wedge \square(\$ \rightarrow \circ^i 1)))].$$

Let us define an equivalence relation on finite traces in  $(0 + 1 + \#)^*$ . For a  $\sigma \in (0 + 1 + \#)^*$ , define  $S(\sigma) = \{w \in (0 + 1)^k \mid \exists \lambda_1, \lambda_2. \lambda_1 \# w \# \lambda_2 = \sigma\}$ . We say that  $\sigma_1 \equiv_k \sigma_2$  if and only if  $S(\sigma_1) = S(\sigma_2)$ . Now observe that the number of equivalence classes of  $\equiv_k$  is  $2^{2^k}$ ; this is because for any  $S \subseteq (0 + 1)^k$ , there is a  $\sigma$  such that  $S(\sigma) = S$ .

Since  $|\phi_k| = \Theta(k^2)$ , it follows that there is some constant  $c'$  such that  $|\phi_k| \leq c'k^2$  for all large enough  $k$ . Let  $c$  be the constant  $1/\sqrt{c'}$ . We will prove this lower bound result by contradiction. Suppose  $\mathcal{A}$  is an LTL forwards monitoring algorithm that uses less than  $2^{c\sqrt{m}}$  space for any LTL formulae of large enough size  $m$ . We will look at the behavior of the algorithm  $\mathcal{A}$  on inputs

of the form  $\phi_k$ . So  $m = |\phi_k| \leq c'k^2$ , and  $\mathcal{A}$  uses less than  $2^k$  space. Since the number of equivalence classes of  $\equiv_k$  is  $2^{2^k}$ , by the pigeon hole principle there must be two strings  $\sigma_1 \not\equiv_k \sigma_2$  such that the memory of  $\mathcal{A}$  on  $\phi_k$  after reading  $\sigma_1\$$  is the same as the memory after reading  $\sigma_2\$$ . In other words,  $\mathcal{A}$  running on  $\phi_k$  will give the same answer on all traces of the form  $\sigma_1\$w$  and  $\sigma_2\$w$ . Now since  $\sigma_1 \not\equiv_k \sigma_2$ , it follows that  $(S(\sigma_1) \setminus S(\sigma_2)) \cup (S(\sigma_2) \setminus S(\sigma_1)) \neq \emptyset$ . Take  $w \in (S(\sigma_1) \setminus S(\sigma_2)) \cup (S(\sigma_2) \setminus S(\sigma_1))$ . Then clearly, exactly one out of  $\sigma_1\$w$  and  $\sigma_2\$w$  is in  $L_k$ , and so  $\mathcal{A}$  running on  $\phi_k$  gives the wrong answer on one of these inputs. Therefore,  $\mathcal{A}$  is not a correct.  $\square$

It seems, however, that this worst-case exponential complexity in the size of the LTL formula is more of theoretical importance than practical, since in general the size of the formula rarely grew more than twice in our experiments. Verification results are very encouraging and show that this optimized semantics is orders of magnitude faster than the first semantics. Traces of less than 10,000 events are verified in milliseconds, while traces of 100,000 events never needed more than 3 seconds. This technique scales quite well; we were able to monitor even traces of hundreds of millions events. As a concrete example, we created an artificial trace by repeating 10 million times the 10 event trace in Subsection 8.1.2, and then checked it against the formula  $\Box(\text{green} \rightarrow \neg \text{red} \cup \text{yellow})$ . There were needed 4.9 billion rewriting steps for a total of about 1,500 seconds. In Subsection 8.3.3 we will see how this algorithm can be made even more efficient, using memoization.

### 8.3.2 Correctness and Completeness

In this subsection we prove that the algorithm presented above is correct and complete with respect to the semantics of finite trace LTL presented in Section 8.1. The proof is done completely in Maude, but since Maude is not intended to be a theorem prover, we actually have to generate the proof obligations by hand. In other words, the proof that follows was *not* generated automatically. However, it could have been mechanized by using proof assistants and/or theorem provers like KUMO [48], PVS [120], or Maude-ITP [22]. We have already done it in PVS, but we prefer to use only plain Maude in this paper.

**Theorem 16** *For any trace  $T$  and any formula  $X$ ,  $T \models X$  if and only if  $T \vdash X$ .*

**Proof:** By induction, both on traces and formulae. We first need to prove two lemmas, namely that the following two equations hold in the context of both LTL and LTL-REVISED:

$$\begin{aligned} (\forall E : \text{Event}, X : \text{Formula}) \quad E \models X &= E \Vdash X, \\ (\forall E : \text{Event}, T : \text{Trace}, X : \text{Formula}) \quad E, T \models X &= T \models X\{E\}. \end{aligned}$$

We prove them by structural induction on the formula  $X$ . Constants  $e$  and  $x$  are needed in order to prove the first lemma via the theorem of constants. However, since we prove these lemmas by structural induction on  $X$ , we not only have to add two constants  $e$  and  $t$  for the universally quantified variables  $E$  and  $T$ , but also two other constants  $y$  and  $z$  standing for formulas which can be combined via operators to give other formulas. The induction hypotheses are added to the following specification via equations. Notice that we merged the two proofs to save space. A proof assistant like KUMO, PVS or Maude-ITP would prove them independently, generating only the needed constants for each of them.

```
fmod PROOF-OF-LEMMAS is
  extending LTL .
  extending LTL-REVISED .
  op e : -> Event .
  op t : -> Trace .
  ops a b c : -> Atom .
  ops y z : -> Formula .
  eq e |= y = e |- y .
  eq e |= z = e |- z .
  eq e,t |= y = t |- y{e} .
  eq e,t |= z = t |- z{e} .
  eq b{e} = true .
  eq c{e} = false .
endfm
```

It is worth reminding the reader at this stage that the functional modules in Maude have initial semantics, so proofs by induction are valid. Before proceeding further, the reader should be aware of the operational semantics of the operation `_==_`, namely that the two argument terms are first reduced to their normal forms which are then compared syntactically (but modulo associativity and commutativity); it returns `true` if and only if the two normal forms are equal. Therefore, the answer `true` means that the two terms are indeed semantically equal, while `false` only means that they could not be proved equal; they can still be equal.

```

reduce (e |= a      == e |- a)
  and (e |= true   == e |- true)
  and (e |= false  == e |- false)
  and (e |= y /\ z == e |- y /\ z)
  and (e |= y ++ z == e |- y ++ z)
  and (e |= [] y   == e |- [] y)
  and (e |= <> y    == e |- <> y)
  and (e |= y U z  == e |- y U z)
  and (e |= o y    == e |- o y)

  and (e,t |= true  == t |= true{e})
  and (e,t |= false == t |= false{e})
  and (e,t |= b     == t |= b{e})
  and (e,t |= c     == t |= c{e})
  and (e,t |= y /\ z == t |= (y /\ z){e})
  and (e,t |= y ++ z == t |= (y ++ z){e})
  and (e,t |= [] y  == t |= ([] y){e})
  and (e,t |= <> y  == t |= (<> y){e})
  and (e,t |= y U z == t |= (y U z){e})
  and (e,t |= o y   == t |= (o y){e}) .

```

It took Maude 129 reductions to prove these lemmas. Therefore, one can safely add now these lemmas as follows:

```

fmod LEMMAS is
  protecting LTL .
  protecting LTL-REVISED .
  var E : Event .
  var T : Trace .
  var X : Formula .
  eq E |= X = E |- X .
  eq E,T |= X = T |= X{E} .
endfm

```

We can now prove the theorem, by induction on traces. More precisely, we show:

$\mathcal{P}(E)$ , and  
 $\mathcal{P}(T)$  implies  $\mathcal{P}(E,T)$ , for all events  $E$  and traces  $T$ ,

where  $\mathcal{P}(T)$  is the predicate “for all formulas  $X$ ,  $T \models X$  iff  $T \vdash X$ ”. This induction schema can be easily formalized in Maude as follows:

```

fmod PROOF-OF-THEOREM is
  protecting LEMMAS .

```

```

    op e : -> Event .
    op t : -> Trace .
    op x : -> Formula .
    var X : Formula .
    eq t |= X = t |- X .
endfm

reduce e   |= x == e   |- x .
reduce e,t |= x == e,t |- x .

```

Notice the difference in role between the constant  $x$  and the variable  $X$ . The first reduction proves the base case of the induction, using the theorem of constants for the universally quantified variable  $X$ . In order to prove the induction step, we first applied the theorem of constants for the universally quantified variables  $E$  and  $T$ , then added  $\mathcal{P}(t)$  to the hypothesis (the equation “ $\text{eq } t \models X = t \vdash X$ .”), and then reduced  $\mathcal{P}(e \ t)$  using again the theorem of constants for the universally quantified variable  $X$ . Like in the proofs of the lemmas, we merged the two proofs to save space.  $\square$

### 8.3.3 Further Optimization by Memoization

Even though the formula transforming algorithm in Subsection 8.3.1 can process 100 million events in about 25 minutes, which is relatively reasonable for practical purposes, it can be significantly improved by adding only 5 more characters to the existing Maude code presented so far. More precisely, one can replace the operation declaration

```
op _{ _ } : Formula Event* -> Formula [prec 10]
```

in module LOGICS-BASIC by the operation declaration

```
op _{ _ } : Formula Event* -> Formula [memo prec 10]
```

The attribute `memo` added to an operation declaration instructs Maude to memorize, or cache, the normal forms of terms rooted in that operation, i.e., those terms will be rewritten only once. Memoization is implemented by hashing, where the entry in the hash table is given by the term to be reduced and the value in the hash is its normal form. In our concrete example, memoization has the effect that any LTL formula will be transformed by a given event exactly once during the monitoring sequence; if the same formula and the same event occur in the future, the resulting modified formula is extracted from the hash table without applying any rewriting step. If one

thinks of LTL in terms of automata, then our new algorithm corresponds to building the monitoring automaton *on the fly*. The obvious benefit of this technique is that only the *needed* part of the automaton is built, namely that part that is reachable during monitoring a particular sequence of events, which is practically very useful because the entire automaton associated to an LTL formula can be exponential in size, so storing it might become a problem.

The use of memoization brings a significant improvement in the case of LTL. For example, the same sequence of 100 million events, which took 1500 seconds using the algorithm presented in Subsection 8.3.1, takes only 185 seconds when one uses memoization, for a total of 2.2 rewritings per processed event and 540,000 events processed per second! We find these numbers amazingly good for any practical purpose we can think of and believe that, taking into account the simplicity, obvious correctness and elegance of the rewriting based algorithm (implemented basically by 8 rewriting rules in LTL-REVISED), it would be hard to argue for any other implementation of LTL monitoring. One should, however, be careful when one uses memoization because hashing slows down the rewriting engine. LTL is a happy case where memoization brings a significant improvement, because the operational semantics of all the operators can be defined recursively, so formulae repeat often during the monitoring process. However, there might be monitoring logics where memoization could be less efficient. Such a logic would probably be an extension of LTL with time, allowing formulae of the form “ $\langle 5 \rangle X$ ” with the meaning “eventually in 5 units of time  $X$ ”, because of a potentially very large number of terms to be memoized:  $\langle 5 \rangle X$ ,  $\langle 4 \rangle X$ , etc. Experimentation is certainly needed if one designs a new logic for monitoring and wants to use memoization.

## 8.4 Generating Forwards, Synchronous and Efficient Monitors

Even though the rewriting based monitoring algorithm presented in the previous section performs quite well in practice, there can be situations in which one wants to minimize the monitoring overhead as much as possible. Additionally, despite its simplicity and elegance, the procedure above requires an efficient rewriting engine modulo associativity and commutativity, which may not be available or may not be desirable on some monitoring platforms, such as, for example, within an embedded system.

In this section we give a technique, based on ideas presented previously in the paper, to generate automata-based optimal monitors for future time LTL formulae. By optimality we here mean everything one may expect, such as minimal number of states, forwards traversal of execution traces, synchronicity, efficiency, but also less standard optimality features, such as transiting from one state to another with a minimum amount of computation. In order to effectively do this we introduce the notion of *binary transition tree* (BTT), as a generalization of binary decision diagrams (BDD) [14], whose purpose is to provide an *optimal order* in which state predicates need to be evaluated to decide the next state. The motivation for this is that in practical applications evaluating a state predicate is a time consuming task, such as for example to check whether a vector is sorted. The associated finite state machines are called *binary transition tree finite state machines* (BTT-FSM).

The drawback of generating an optimal BTT-FSM statically, i.e., before monitoring, is the worst-case double exponential time/space required at startup. Therefore, the algorithm presented in this section is recommended for situations where the LTL formulae to monitor are relatively small in size but the runtime overhead is desired to be minimal. It is worth noting that the BTT-FSM generation process can potentially take place on a machine different from the one performing the monitoring. In particular, one can think of a WWW fast server offering LTL-to-BTT-FSM services via the Internet, which can also maintain a database of already generated BTT-FSMs to avoid regenerating the same monitors.

#### 8.4.1 From LTL Formulae to BTT-FSMs

Informally, our algorithm to generate optimal BTT-FSMs from LTL formulae consists of two steps. First, it generates an MT-FSM with a minimum number of states. Then, using the technique presented in the previous subsection, it generates an optimal BTT from each MT.

To generate a minimal MT-FSM, our algorithm uses the rewriting based procedure presented in Section 8.3 on all possible events, until the set of formulae to which the original LTL formula can “evolve” stabilizes. The procedure LTL2MT-FSM shown in Figure 8.1 builds an MT-FSM whose states are formulae, with the help of a validity checker. Initially, the set  $S$  of states contains only the original LTL formula. LTL2MT-FSM is called on the original LTL formula, and then recursively in a depth-first manner on all the formulae to which the initial formula can ever evolve via the

```

1. let  $S$  be  $\varphi$ 
2. procedure LTL2MT-FSM( $\varphi$ )
3.   let  $\mu^*(\varphi)$  be  $\emptyset$ 
4.   let  $\mu(\varphi)$  be  $\emptyset$ 
5.   foreach  $\theta : A \rightarrow \{true, false\}$  do
6.     let  $e_\theta$  be the list of atoms  $a$  with  $\theta(a) = true$ 
7.     let  $p_\theta$  be the proposition  $\bigwedge\{a \mid \theta(a) = true\} \wedge \bigwedge\{\neg a \mid \theta(a) = false\}$ 
8.     let  $\mu^*(\varphi)$  be MERGE( $[p_\theta ? \varphi\{e_\theta^*\}], \mu^*(\varphi)$ )
9.     let  $\varphi_\theta$  be  $\varphi\{e_\theta\}$ 
10.    if there is  $\varphi' \in S$  with VALID( $\varphi_\theta \leftrightarrow \varphi'$ )
11.      then let  $\mu(\varphi)$  be MERGE( $[p_\theta ? \varphi']$ ,  $\mu(\varphi)$ )
12.    else let  $S$  be  $S \cup \{\varphi_\theta\}$ 
13.      let  $\mu(\varphi)$  be MERGE( $[p_\theta ? \varphi_\theta]$ ,  $\mu(\varphi)$ )
14.      LTL2MT-FSM( $\varphi_\theta$ )
15.    endfor
16.    if  $\mu(\varphi) = [true ? \varphi]$  and  $\mu^*(\varphi) = [true ? b]$  then replace  $\varphi$  by  $b$  everywhere
17. endprocedure

```

Figure 8.1: Algorithm to generate a minimal MT-FSM\*  $(S, A, \mu, \mu^*, \varphi)$  from an LTL formula  $\varphi$ .

event-consuming operator  $_{-}\{\_ \}$  introduced in Section 8.3.

For each LTL state formula  $\varphi$  in  $S$ , multi-transitions  $\mu(\varphi)$  and  $\mu^*(\varphi)$  are maintained. For each possible event  $\theta$ , one first updates  $\mu^*(\varphi)$  by considering the case in which  $\theta$  is the last event in a trace (step 8), and then the current formula  $\varphi$  evolves into the corresponding formula  $\varphi_\theta$  (step 9). If some equivalent formula  $\varphi'$  to  $\varphi_\theta$  has already been discovered then one only needs to modify the multi-transition set of  $\varphi$  accordingly in order to point to  $\varphi'$  (step 11). Notice that equivalence of LTL formulae is checked by using a validity procedure (step 10), which is given in Figure 8.2. If there is no formula in  $S$  equivalent to  $\varphi_\theta$ , then the new formula  $\varphi_\theta$  is added to  $S$ , multi-transition  $\mu(\varphi)$  is updated accordingly, and then the MT-FSM generation procedure is called recursively on  $\varphi_\theta$ . This way, one eventually generates all possible LTL formulae into which the initial formula can ever evolve during a monitoring session; this happens, of course, modulo finite trace LTL semantic equivalence, implemented elegantly using the validity checker

described below. By Theorem 15, this recursion will eventually terminate, leading to an MT-FSM\*.

The procedure `VALID` used in LTL2MT-FSM above is defined in Figure 8.2. It essentially follows the same idea of generating all possible formulae to which the original LTL formula tested for validity can evolve via the event consumption operator defined by rewriting in Section 8.3, but for each newly obtained formula  $\varphi$  and for each event  $\theta$ , it also checks whether an execution trace stopping at that event would be a rejecting sequence. The intuition for this check is that a formula is valid under finite trace LTL semantics if and only if any (finite) sequence of events satisfies that formula; since any generated formula corresponds to one into which the initial formula can evolve, we need to make sure that each of these formulae becomes *true* under any possible last monitored event. This is done by rewriting the term  $\varphi\{e_\theta^*\}$  with the rules in Section 8.3 to its normal form (step 5.), i.e., *true* or *false*. The formula is valid if and only if there is no rejecting sequence; the entire space of evolving formulae is again explored by depth-first search. Notice

```

1. let  $S$  be  $\varphi$ 
2. function VALID( $\varphi$ )
3.   foreach  $\theta : A \rightarrow \{true, false\}$  do
4.     let  $e_\theta$  be the list of atoms  $a$  with  $\theta(a) = true$ 
5.     if  $\varphi\{e_\theta^*\} = false$  then return false
6.     let  $\varphi_\theta$  be  $\varphi\{e_\theta\}$ 
7.     if  $\varphi_\theta \notin S$ 
8.       then let  $S$  be  $S \cup \{\varphi_\theta\}$ 
9.       if VALID( $\varphi_\theta$ ) = false then return false
10.    endfor
11.    return true
12. end function

```

Figure 8.2: Validity checker for an LTL formula  $\varphi$ .

that `VALID` does not test for equivalence of formulae, so it can potentially generate a larger number of formulae than LTL2MT-FSM. However, by Theorem 15, this procedure will also eventually terminate.

**Theorem 17** *Given an LTL formula  $\varphi$  of size  $m$ , the following hold:*

1. The procedure `VALID` is correct, that is, `VALID( $\varphi$ )` returns true if and only if  $\varphi$  is satisfied, as defined in Section 8.1, by any finite trace;
2. The space and time complexity of `VALID( $\varphi$ )` is  $2^{O(2^m)}$ ;

Additionally, letting  $M$  denote the  $MT\text{-}FSM^*$  ( $S, A, \mu, \mu^*, \varphi$ ) generated by `LTL2MT-FSM( $\varphi$ )`, the following hold:

3. `LTL2MT-FSM( $\varphi$ )` is correct; more precisely,  $M$  has the property that for any events  $\theta_1, \dots, \theta_n, \theta$ , it is the case that  $\varphi \xrightarrow{\theta_1} \varphi_1 \dots \xrightarrow{\theta_n} \varphi_n \xrightarrow{\theta^*} \text{true}$  if and only if the finite trace  $e_{\theta_1} \dots e_{\theta_n} e_{\theta}$  satisfies  $\varphi$ , and  $\varphi \xrightarrow{\theta_1} \varphi_1 \dots \xrightarrow{\theta_n} \varphi_n \xrightarrow{\theta^*} \text{false}$  if and only if  $e_{\theta_1} \dots e_{\theta_n} e_{\theta}$  does not satisfy  $\varphi$ ;
4.  $M$  is synchronous; more precisely, for any events  $\theta_1, \dots, \theta_n$ , it is the case that  $\varphi \xrightarrow{\theta_1} \varphi_1 \dots \xrightarrow{\theta_n} \varphi_n \xrightarrow{e_{\theta}} \text{true}$  if and only if  $e_{\theta_1} \dots e_{\theta_n} e_{\theta}$  is a valid prefix of  $\varphi$ , and  $\varphi \xrightarrow{\theta_1} \varphi_1 \dots \xrightarrow{\theta_n} \varphi_n \xrightarrow{e_{\theta}} \text{false}$  if and only if  $e_{\theta_1} \dots e_{\theta_n} e_{\theta}$  is a bad prefix of  $\varphi$ ;
5. The space and time complexity of `LTL2MT-FSM( $\varphi$ )` is  $2^{O(2^m)}$ ;
6.  $M$  is the smallest  $MT\text{-}FSM^*$  which is correct and synchronous (as defined in 3 and 4 above);
7. Monitoring against a  $BTT\text{-}FSM^*$  corresponding to  $M$ , as shown in Subsection 6.1.2, needs  $O(2^m)$  space and time.

**Proof:** 1. Using a depth-first search strategy, the procedure `VALID` visits all possible formulae into which the original formula  $\varphi$  can evolve via any possible sequence of events. These formulae are different modulo the rewriting rules in Section 8.3 which are used to simplify LTL formulae and to remove the derivatives, but those rules were shown to be sound, so `VALID` indeed explores all possible LTL formulae into which  $\varphi$  can *semantically* evolve. Moreover, for each such formula, say  $\varphi'$ , and each event, `VALID` checks at Step 5 whether a trace terminating with that event in  $\varphi'$  would lead to rejection. `VALID( $\varphi$ )` returns true if and only if there is no such rejection, so it is indeed correct: if it had been some finite trace that did not satisfy  $\varphi$  then `VALID` would have found it during its exhaustive search.

2. The space required by `VALID( $\varphi$ )` is clearly dominated by the size of  $S$ . By Theorem 15, each formula  $\varphi'$  into which  $\varphi$  can evolve needs  $O(2^m)$  space to be stored. That means that there can be at most  $2^{O(2^m)}$  such formulae. So

the total space needed to store  $S$  in the worst case is  $2^{O(2^m)}$ . An amortized analysis of its running time, tells us that `VALID` runs its “for each event” loop one per formula in  $S$ . Since the number of events is much less than  $2^m$  and since reducing the derivative of a formula by an event takes time proportional with the size of the formula, we deduce that the total running time of the `VALID` procedure is  $2^{O(2^m)}$ .

3. By Theorem 16 and the event consuming procedure described in Subsection 8.3.1, it follows that  $e_{\theta_1} \dots e_{\theta_n} e_\theta$  satisfies  $\varphi$  if and only if  $\varphi\{e_{\theta_1}\}\{\dots\}\{e_{\theta_n}\}\{e_\theta^*\}$  reduces to *true*, and that  $e_{\theta_1} \dots e_{\theta_n} e_\theta$  does not satisfy  $\varphi$  if and only if  $\varphi\{e_{\theta_1}\}\{\dots\}\{e_{\theta_n}\}\{e_\theta^*\}$  reduces to *false*. It is easy to see that `VALID`( $\psi \leftrightarrow \psi'$ ) returns true if and only if  $\psi$  and  $\psi'$  are formulae equivalent under the finite trace LTL semantics. That means the MT-FSM\* generated by `LTL2MT-FSM`( $\varphi$ ) contains a formula-state  $\varphi_1$  which is equivalent to  $\varphi\{e_{\theta_1}\}$ ; moreover,  $\varphi \xrightarrow{\theta_1} \varphi_1$  in this MT-FSM\*. Inductively, one can show that there is a series of formulae-states  $\varphi_1, \dots, \varphi_n$  such that  $\varphi \xrightarrow{\theta_1} \varphi_1 \dots \xrightarrow{\theta_n} \varphi_n$  is a sequence of transitions in the generated MT-FSM and  $\varphi\{e_{\theta_1}\}\{\dots\}\{e_{\theta_n}\}$  is equivalent to  $\varphi_n$ . The rest follows by noting that  $\varphi_n\{e_\theta^*\}$  reduces to true if and only if  $\varphi_n \xrightarrow{\theta^*} \text{true}$  in the generated MT-FSM, and that  $\varphi_n\{e_\theta^*\}$  reduces to false if and only if  $\varphi_n \xrightarrow{\theta^*} \text{false}$  in the MT-FSM.

4. As in 3 above, one can inductively show that there is a series of formulae  $\varphi_1, \dots, \varphi_n, \varphi'$  such that  $\varphi \xrightarrow{\theta_1} \varphi_1 \dots \xrightarrow{\theta_n} \varphi_n \xrightarrow{\theta} \varphi'$  is a sequence of transitions in the generated MT-FSM and  $\varphi\{e_{\theta_1}\}\{\dots\}\{e_{\theta_n}\}\{e_\theta\}$  is equivalent to  $\varphi'$ . By Proposition 14, due to the use of the validity checker in step 10 of `LTL2MT-FSM` it follows that  $e_{\theta_1} \dots e_{\theta_n} e_\theta$  is a valid prefix of  $\varphi$  if and only if  $\varphi'$  is equivalent to *true*, and that  $e_{\theta_1} \dots e_{\theta_n} e_\theta$  is a bad prefix of  $\varphi$  if and only if  $\varphi'$  is equivalent to *false*. The rest follows now by Step 16 in the algorithm in Figure 8.1, which, due to the validity checker, provides a necessary and sufficient condition for a processed formula to be semantically equivalent to *true* or *false*, respectively.

5. The space required by `LTL2MT-FSM`( $\varphi$ ) is again dominated by the size of  $S$ . Like in the analysis of `VALID` in 2 above, by Theorem 15 we get that there can be at most  $2^{O(2^m)}$  formulae generated in  $S$ , so the total space needed to store  $S$  in the worst case is also  $2^{O(2^m)}$ . For each newly added formula in  $S$  and for each event, Step 10 calls the procedure `VALID` potentially once for each already existing formula in  $S$ . It is important to

notice that the formulae on which VALID is called are exclusive disjunctions of conjunctions of sub-formulae rooted in temporal operators of the original formula  $\varphi$ , so its space and time complexity will also be  $2^{O(2^m)}$  each time it is called by LTL2MT-FSM. One can easily see now that the space and time requirements of LTL2MT-FSM( $\varphi$ ) are also  $2^{O(2^m)}$  (the constant in the exponent  $O(2^m)$  can be appropriately enlarged).

6. For any MT-FSM\* machine  $M' = (S', A, \mu', \mu'^*, q_0)$ , one can associate a formula, more precisely a state in  $M$ , to each state in  $S'$  as follows.  $\varphi$  is associated to the initial state  $q_0$ . Then for each transition  $q_1 \xrightarrow{\theta} q_2$  in  $M'$  such that  $q_1$  has a formula  $\varphi_1$  associated and  $q_2$  has no formula associated, then one associates that  $\varphi_2$  to  $q_2$  with the property that  $\varphi_1 \xrightarrow{\theta} \varphi_2$  is a transition in  $M$ . For a state  $q$  in  $S'$  let  $\varphi_q$  be the formula in  $S$  associated to it. Since  $M'$  is correct and synchronous for  $\varphi$ , it follows that  $\mathcal{L}_{M'}(q) = \mathcal{L}_M(\varphi_q)$  for any state  $q$  in  $S'$ . We can now show that the map associating a formula in  $S$  to any state in  $S'$  is surjective, which shows that  $M$  has therefore a minimal number of states. Let  $\varphi'$  be a formula in  $S$  and let  $\varphi \xrightarrow{\theta_1} \varphi_1 \cdots \xrightarrow{\theta_n} \varphi_n \xrightarrow{\theta} \varphi'$  be a sequence of transitions in  $M$  leading to  $\varphi'$ ; note that all formulae in  $S$  are reachable via transitions in  $M$  from  $\varphi$ . Let  $q'$  be the state in  $S'$  such that  $q_0 \xrightarrow{\theta_1} q_1 \cdots \xrightarrow{\theta_n} q_n \xrightarrow{\theta} q'$ . Then  $\mathcal{L}_{M'}(q') = \mathcal{L}_M(\varphi_{q'})$ . Since by Proposition 14,  $\mathcal{L}_{M'}(q') = \{t \mid e_{\theta_1} \dots e_{\theta_n} e_{\theta} t \in \mathcal{L}_{M'}(q_0)\}$  and  $\mathcal{L}_M(\varphi') = \{t \mid e_{\theta_1} \dots e_{\theta_n} e_{\theta} t \in \mathcal{L}_M(\varphi)\}$ , it follows that  $\mathcal{L}_M(\varphi_{q'}) = \mathcal{L}_M(\varphi')$ , that is, that  $\varphi_{q'}$  and  $\varphi'$  are equivalent. Since Step 10 of the LTL2MT-FSM eventually uses the validity checker on any pairs of formulae in  $S$ , it follows that  $\varphi_{q'} = \varphi'$ .

7. In order to distinguish  $N$  pieces of data,  $\log(N)$  bits are needed to encode each datum. Therefore, one needs  $O(2^m)$  bits to encode a state of  $M$ , which is the main memory needed by the monitoring algorithm. Like in Theorem 15, we assume “idealistic” rewriting engines able to optimize the space requirements; we are not aware whether Maude is able to attain this performance or not. To make a transition from one state to another, a BTT-FSM\* associated to the MT-FSM\* generated by LTL2MT-FSM( $\varphi$ ) needs to only evaluate at most all the atomic predicates occurring in the formula, which, assuming that evaluating atom predicates takes unit time, is clearly  $O(2^m)$ . When the entire BTT is evaluated, it finally has to store the newly obtained state of the BTT-FSM\*, which can reuse the space of the previous one,  $O(2^m)$ , and which also takes time  $O(2^m)$ .  $\square$

Once the procedure LTL2MT-FSM terminates, the formulae  $\varphi$ ,  $\varphi'$ , etc.,

are not needed anymore, so one can and should replace them by unique labels in order to reduce the amount of storage needed to encode the MT-FSM\* and/or the corresponding BTT-FSM\*. This algorithm can be relatively easily implemented in any programming language. We have, however, found Maude again a very elegant system for this task, implementing the entire LTL formula to BTT-FSM algorithm in about 200 lines of code.

### 8.4.2 Examples

The BTT-FSM generation algorithm presented in this section, despite its overall worst-case high startup time, can be very useful when formulae are relatively short, as it is most often the case in practice. For the traffic light controller requirement formula discussed previously in the paper,  $\Box(\text{green} \rightarrow (\neg \text{red}) \cup \text{yellow})$ , this algorithm generates in about 0.2 seconds the optimal BTT-FSM\* in Figure 8.3, also shown in Figure 6.3 in flowchart notation; Figure 6.2 shows its optimal MT-FSM\*. For simplicity, the states *true* and

State	Non-terminal event	Terminal event
1	<code>yellow ? 1 : green ? red ? false : 2 : 1</code>	<code>yellow ? true : green ? false : true</code>
2	<code>yellow ? 1 : red ? false : 2</code>	<code>yellow ? true : false</code>

Figure 8.3: An optimal BTT-FSM\* for the formula  $\Box(\text{green} \rightarrow \neg \text{red} \cup \text{yellow})$ .

*false* do not appear in Figure 8.3. Notice that the atomic predicate `red` does *not* need to be evaluated on terminal events and that `green` does not need to be evaluated in state 2. In this example, the colors are not supposed to exclude each other, that is, the traffic controller can potentially be both green and red.

The LTL formulae on which our algorithm has the worst performance are those containing many nested temporal operators (which are not frequently used in specifications anyway, because of the high risk of getting them wrong). For example, it takes our Maude implementation of this algorithm 1.3 seconds to generate the minimal 3-state (*true* and *false* states are not counted) BTT-FSM\* for the formula  $a \cup (b \cup (c \cup a))$  and 13.2 seconds to generate the 7-state minimal BTT-FSM\* for the formula  $((a \cup b) \cup c) \cup a$ . It never took our current implementation more than a few seconds to generate the BTT-FSM\* of any LTL formula of interest for our applications, i.e.,

non-artificial. Figure 8.4 shows the generated BTT-FSM of some artificial LTL formulae, taking together less than 15 seconds to be generated. To keep the figure small, the states *true* and *false* together with their self-transitions are not shown in Figure 8.4, and they are replaced by **t** and **f** in BTTs.

Formula	State	Monitoring BTT	Terminating BTT
$\Box \diamond a$	1	1	$a ? \mathbf{t} : \mathbf{f}$
$\diamond(\Box a \vee \Box \neg a)$	—	—	—
$\Box(a \rightarrow \diamond b)$	1	$a ? (b ? 1 : 2) : 1$	$a ? (b ? \mathbf{t} : \mathbf{f}) : \mathbf{t}$
	2	$b ? 1 : 2$	$b ? \mathbf{t} : \mathbf{f}$
$a \mathcal{U} (b \mathcal{U} c)$	1	$c ? \mathbf{t} : (a ? 1 : (b ? 2 : \mathbf{f}))$	$c ? \mathbf{t} : \mathbf{f}$
	2	$c ? \mathbf{t} : (b ? 2 : \mathbf{f})$	$c ? \mathbf{t} : \mathbf{f}$
$a \mathcal{U} (b \mathcal{U} (c \mathcal{U} d))$	1	$d ? \mathbf{t} : a ? 1 : b ? 2 : c ? 3 : \mathbf{f}$	$d ? \mathbf{t} : \mathbf{f}$
	2	$d ? \mathbf{t} : b ? 2 : c ? 3 : \mathbf{f}$	$d ? \mathbf{t} : \mathbf{f}$
	3	$d ? \mathbf{t} : c ? 3 : \mathbf{f}$	$d ? \mathbf{t} : \mathbf{f}$
$((a \mathcal{U} b) \mathcal{U} c) \mathcal{U} d$	1	$d ? \mathbf{t} : c ? 1 : b ? 4 : a ? 5 : \mathbf{f}$	$d ? \mathbf{t} : \mathbf{f}$
	2	$b ? c ? \mathbf{t} : 7 : a ? c ? 6 : 2 : \mathbf{f}$	$c ? b ? \mathbf{t} : \mathbf{f} : \mathbf{f}$
	3	$b ? d ? \mathbf{t} : c ? 1 : 4 : a ? d ? 6 : c ? 3 : 5 : \mathbf{f}$	$d ? b ? \mathbf{t} : \mathbf{f} : \mathbf{f}$
	4	$c ? d ? \mathbf{t} : 1 : b ? d ? 7 : 4 : a ? d ? 2 : 5 : \mathbf{f}$	$d ? c ? \mathbf{t} : \mathbf{f} : \mathbf{f}$
	5	$b ? d ? c ? \mathbf{t} : 7 : c ? 1 : 4 : a ? d ? c ? 6 : 2 : c ? 3 : 5 : \mathbf{f}$	$d ? c ? b ? \mathbf{t} : \mathbf{f} : \mathbf{f} : \mathbf{f}$
	6	$b ? \mathbf{t} : a ? 6 : \mathbf{f}$	$b ? \mathbf{t} : \mathbf{f}$
	7	$c ? \mathbf{t} : b ? 7 : a ? 2 : \mathbf{f}$	$c ? \mathbf{t} : \mathbf{f}$

Figure 8.4: Six BTT-FSM\*s generated in less than 15 seconds.

The generated BTT-FSM\*s are monitored most efficiently on RAM machines, due to the fact that conditional statements are implemented via jumps in memory. Monitoring BTT-FSM\*s using rewriting does not seem appropriate because it would require linear time, as a function of number of states, to extract the BTT associated to a state in a BTT-FSM\*. However, we believe that the algorithm presented in Section 8.3 is satisfactory in practice if one is willing to use a rewriting engine for monitoring.

## 8.5 Conclusions

This paper presented a foundational study in using rewriting in runtime verification and monitoring of systems. After a short discussion on types of monitoring and mathematical and technological preliminaries, a finite trace linear temporal logic was defined, together with an immediate but inefficient implementation of a monitor following directly its semantics. Then an

efficient but ineffective implementation based on dynamic programming was presented, which traverses the execution trace backwards. The first effective and relatively efficient rewriting algorithm was further introduced, based on the idea of transforming the monitoring requirements as events are received from the monitored program. A non-trivial improvement of this algorithm based on hashing rewriting results, thereby reducing the number of rewritings performed during trace analysis, was also proposed. The hashing corresponds to building an observer automaton on-the-fly, having the advantage that only the part of the automaton that is needed for analyzing a given trace is generated. The resulting algorithm is very efficient. Since in many cases one would want to generate an observer finite state machine (or automaton) a priori, for example when a rewriting system cannot be used for monitoring, or when minimal runtime overhead is needed, a specialized data-structure called a binary transition tree (BTT) and corresponding finite state machines were introduced, and an algorithm for generating minimal such monitors from temporal formulae was discussed.

All algorithms were implemented in surprisingly few lines of Maude code, illustrating the strength of rewriting for this particular domain. In spite of the reduced size of the code, the implementations seem to be efficient for practical purposes. As a consequence, we have demonstrated how rewriting can be used not only to experiment with runtime monitoring logics, but also as an implementation language. As an example of future work is the extension of LTL with real-time constraints. Since Maude by itself provides a high-level specification language, one can argue that Maude in its entirety can be used for writing requirements. Further work will show whether this avenue is fruitful. Some of the discussed results and algorithms have been already used in two NASA applications, JPAX and X9, but an extensive experimental assesment of these techniques is left as future work.

## Chapter 9

# Efficient Monitoring of $\omega$ -Languages

---

*currently from CAV'05 paper [30]; this material will eventually be dissolved in other parts of the book.*

---

**Abstract:** We present a technique for generating efficient monitors for  $\omega$ -regular-languages. We show how Büchi automata can be reduced in size and transformed into special, statistically optimal nondeterministic finite state machines, called *binary transition tree finite state machines (BTT-FSMs)*, which recognize precisely the minimal bad prefixes of the original  $\omega$ -regular-language. The presented technique is implemented as part of a larger monitoring framework and is available for download.

### 9.1 Introduction

There is increasing recent interest in the area of *runtime verification* [65, 123], which is an area which aims at bridging testing and formal verification. In runtime verification, monitors are generated from system requirements. These monitors observe online executions of programs and check them against requirements. The checks can be either *precise*, with the purpose of detecting existing errors in the observed execution trace, or *predictive*, with the purpose of detecting errors that have not occurred in the observed execution but were “close to happening” and could possibly occur in other executions of the (typically concurrent) system. Runtime verification can be used either during

testing, to catch errors, or during operation, to detect and recover from errors. Since monitoring unavoidably adds runtime overhead to a monitored program, an important technical challenge in runtime verification is that of synthesizing *efficient* monitors from specifications.

Requirements of systems can be expressed in a variety of formalisms, not all of them necessarily easily monitorable. As perhaps best shown by the immense success of programming languages like Perl and Python, regular patterns can be easily devised and understood by ordinary software developers.  $\omega$ -regular-languages [15, 129] add infinite repetitions to regular languages, thus allowing one to specify properties of reactive systems [90]. The usual acceptance condition in finite state machines (FSM) needs to be modified in order to recognize infinite words, thus leading to Büchi automata [21]. Logics like linear temporal logics (LTL) [90] often provide a more intuitive and compact means to specify system requirements than  $\omega$ -regular patterns. It is therefore not surprising that a large amount of work has been dedicated to generating (small) Büchi automata from, and verifying programs against, LTL formulae [43, 129, 34, 41].

Based on the belief that  $\omega$ -languages represent a powerful and convenient formalism to express requirements of systems, we address the problem of generating efficient monitors from  $\omega$ -languages expressed as Büchi automata. More precisely, we generate monitors that recognize the *minimal bad prefixes* [83] of such languages. A bad prefix is a *finite* sequence of events which cannot be the prefix of any accepting trace. A bad prefix is minimal if it does not contain any other bad prefix. Therefore, our goal is to develop efficient techniques that read events of the monitored program incrementally, and precisely detect when a *bad prefix* has occurred. Dual to the notion of bad prefix is that of a good prefix, meaning that the trace will be accepted for any infinite extension of the prefix.

We present a technique that transforms a Büchi automaton into a special (nondeterministic) finite state machine, called a *binary transition tree finite state machine (BTT-FSM)*, that can be used as a monitor: by maintaining a set of possible states which is updated as events are available. A sequence of events is a bad prefix iff the set of states in the monitor becomes empty. One interesting aspect of the generated monitors is that they may contain a special state, called *neverViolate*, which, once reached, indicates that the specification is *not monitorable* from that moment on. That can mean either that the specification has been fulfilled (e.g., a specification  $\diamond(x > 0)$  becomes fulfilled when  $x$  is first seen larger than 0), or that from that moment

on, there will always be some possible continuation of the execution trace. For example, the monitor generated for  $\Box(a \rightarrow \Diamond b)$  will have exactly one state, *neverViolate*, reflecting the intuition that liveness properties cannot be monitored.

As usual, a program state is abstracted as a set of relevant atomic predicates that hold in that state. However, in the context of monitoring, the evaluation of these atomic predicates can be the most expensive part of the entire monitoring process. One predicate, for example, can say whether the vector  $v[1\dots 1000]$  is sorted. Assuming that each atomic predicate has a given evaluation cost and a given probability to hold, which can be estimated apriori either by static or by dynamic analysis, the BTT-FSM generated from a Büchi automaton executes a “conditional program”, called a *binary transition tree (BTT)*, evaluating atomic predicates *by need* in each state in order to statistically optimize the decision to which states to transit. One such BTT is shown in Fig. 9.2.

The work presented in this paper is part of a larger project focusing on *monitoring-oriented programming (MOP)* [19, 18] which is a tool-supported software development framework in which monitoring plays a foundational role. MOP aims at reducing the gap between specification and implementation by integrating the two through monitoring: specifications are checked against implementations at runtime, and recovery code is provided to be executed when specifications are violated. MOP is specification-formalism-independent: one can add one’s favorite or domain-specific requirements formalism via a generic notion of *logic plug-in*, which encapsulates a formal logical syntax plus a corresponding monitor synthesis algorithm. The work presented in this paper is implemented and provided as part of the LTL logic plugin in our MOP framework. It is also available for online evaluation and download on the MOP website [1].

**Some Background and Related Work.** Automata theoretic model-checking is a major application of Büchi automata. Many model-checkers, including most notably SPIN [67], use this technique. So a significant effort has been put into the construction of small Büchi automata from LTL formulae. Gerth *et al.* [43] show a tableaux procedure to generate on-the-fly Büchi automata of size  $2^{O(|\varphi|)}$  from LTL formulae  $\varphi$ . Kesten *et al.* [75] describe a backtracking algorithm, also based on tableaux, to generate Büchi automata from formulae involving both past and future modalities (PTL), but no complexity results are shown. It is known that LTL model-checking is PSPACE-complete [122] and PTL is as expressive and as hard as LTL [91],

though exponentially more succinct [91]. Recently, Gastin and Oddoux [41] showed a procedure to generate standard Büchi automata of size  $2^{O(|\varphi|)}$  from PTL via alternating automata. Several works [34, 43] describe simplifications to reduce the size of Büchi automata. Algebraic simplifications can also be applied *apriori* on the LTL formula. For instance,  $a \mathcal{U} b \wedge c \mathcal{U} b \equiv (a \wedge c) \mathcal{U} b$  is a valid LTL congruence that will reduce the size of the generated Büchi automaton. All these techniques producing small automata are very useful in our monitoring context because the smaller the original Büchi automaton for the  $\omega$ -language, the smaller the BTT-FSM. Simplifications of the automaton *with respect to monitoring* are the central subject of this paper.

Kupferman *et al.* [83] classify safety according to the notion of *informativeness*. Informative prefixes are those that “tell the whole story”: they witness the violation (or validation) of a specification. Unfortunately, not all bad prefixes are informative; e.g., the language denoted by  $\Box(q \vee \circ(\Box(p))) \wedge \Box(r \vee \circ(\Box(\neg p)))$  does not include any word whose prefix is  $\{q, r\}, \{q\}, \{!p\}$ . This is a (minimal) bad but not informative prefix, since it does not witness the violation taking place in the next state. One can use the construction described in [83] to build an automaton of size  $O(2^{2^{|\varphi|}})$  which recognizes all bad prefixes but, unfortunately, this automaton may be too large to be stored. Our fundamental construction is similar in spirit to theirs but we do not need to apply a subset construction on the input Büchi since we already maintain the set of possible states that the running program can be in. Geilen [42] shows how Büchi automata can be turned into monitors. The construction builds a tableaux similar to [43] in order to produce an FSM of size  $O(2^{|\varphi|})$  for recognizing informative good prefixes. Here we detect all the minimal bad prefixes, rather than just the informative ones. Unlike model-checking where a user hopes to see a counter-example that witnesses the violation, when monitoring critical applications one might want to observe a problem as soon as it occurs.

The technique illustrated here is implemented as a plug-in in the MOP *runtime verification (RV)* framework [19, 18]. Other RV tools include JAVA-MAC [76], JPAX [56], JMPAX [119], and EAGLE [12]. JAVA-MAC uses a special interval temporal logic as the specification language, while JPAX and JMPAX support variants of LTL. These systems instrument the JAVA bytecode to emit events to an external monitor observer. JPAX was used to analyze NASA’s K9 Mars Rover code [10]. JMPAX extends JPAX with predictive capabilities. EAGLE is a finite-trace temporal logic and tool for runtime verification, defining a logic similar to the  $\mu$ -calculus with data-

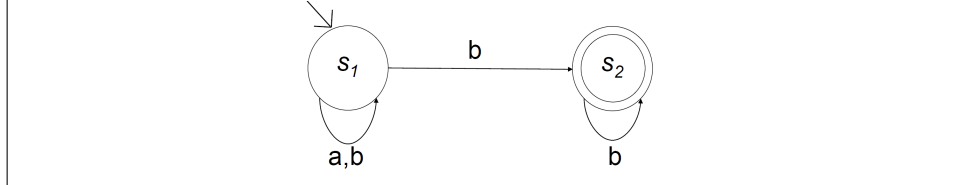


Figure 9.1: Büchi automaton recognizing the  $\omega$ -regular expression  $(a + b)^*b^\omega$

parameterization

## 9.2 Preliminaries: Büchi Automata

Büchi automata and their  $\omega$ -languages have been studied extensively during the past decades. They are well suited to program verification because one can check satisfaction of properties represented as Büchi automata statically against transition systems [129, 21]. LTL is an important but proper subset of  $\omega$ -languages.

**Definition 31** A (nondeterministic) standard **Büchi automaton** is a tuple  $\langle \Sigma, S, \delta, S_0, \mathcal{F} \rangle$ , where  $\Sigma$  is an **alphabet**,  $S$  is a set of **states**,  $\delta: S \times \Sigma \rightarrow 2^S$  is a **transition function**,  $S_0 \subseteq S$  is the set of **initial states**, and  $\mathcal{F} \subseteq S$  is a set of **accepting states**.

In practice,  $\Sigma$  typically refers to events or actions in a system to be analyzed.

**Definition 32** A Büchi automaton  $\mathcal{A} = \langle \Sigma, S, \delta, S_0, \mathcal{F} \rangle$  is said to **accept** an infinite word  $\tau \in \Sigma^\omega$  iff there is some **accepting run** in the automaton, that is, a map  $\rho: \mathbb{N} \rightarrow S$  such that  $\rho_0 \in S_0$ ,  $\rho_{i+1} \in \delta(\rho_i, \tau_i)$  for all  $i \geq 0$ , and  $\text{inf}(\rho) \cap \mathcal{F} \neq \emptyset$ , where  $\text{inf}(\rho)$  contains the states occurring infinitely often in  $\rho$ . The **language of  $\mathcal{A}$** ,  $\mathcal{L}(\mathcal{A})$ , consists of all words it accepts.

Therefore,  $\rho$  can be regarded as an infinite path in the automaton that starts with an initial state and contains at least one accepting state appearing infinitely often in the trace. Fig. 9.1 shows a nondeterministic Büchi automaton for the  $\omega$ -regular expression  $(a + b)^*b^\omega$  that contains all the infinite words over  $a$  and  $b$  with finitely many  $a$ s.

**Definition 33** Let  $\mathcal{L}(\mathcal{A})$  be the language of a Büchi automaton  $\mathcal{A} = \langle \Sigma, S, \delta, S_0, \mathcal{F} \rangle$ . A finite word  $x \in \Sigma^*$  is a **bad prefix of  $\mathcal{A}$**  iff for any  $y \in \Sigma^\omega$  the concatenation  $xy \notin \mathcal{L}(\mathcal{A})$ . A bad prefix is **minimal** if no other bad prefix is a prefix of it.

Therefore, no bad prefix of the language of a Büchi automaton can be extended to an accepted word. Similarly to [21], from now on we may tacitly assume that  $\Sigma$  is defined in terms of propositions over atoms. For instance, the self-transitions of  $s_1$  in Fig. 9.1 can be represented as one self-transition,  $a \vee b$ .

### 9.3 Multi-Transitions and Binary Transition Trees

Büchi automata cannot be used unchanged as monitors. For the rest of the paper we explore structures suitable for monitoring as well as techniques to transform Büchi automata into such structures. Deterministic multi-transitions (*MT*) and binary-transition trees (*BTTs*) were introduced in [60, 106]. In this section we extend their original definitions with nondeterminism.

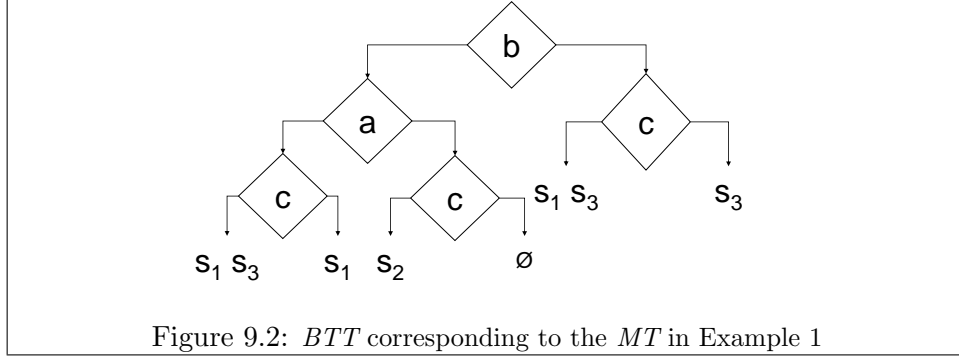
**Definition 34** *Let  $S$  and  $A$  be sets of **states** and **atomic predicates**, respectively, and let  $P_A$  denote the set of **propositions** over atoms in  $A$ , using the usual boolean operators. If  $\{s_1, s_2, \dots, s_n\} \subseteq S$  and  $\{p_1, p_2, \dots, p_n\} \subseteq P_A$ , we call the  $n$ -tuple  $[p_1: s_1, p_2: s_2, \dots, p_n: s_n]$  a **(nondeterministic) multi-transition (MT)** over  $P_A$  and  $S$ . Let  $MT(P_A, S)$  denote the set of MTs over  $P_A$  and  $S$ .*

Intuitively, if a monitor is in a state associated to an *MT*  $[p_1: s_1, p_2: s_2, \dots, p_n: s_n]$  then  $p_1, p_2, \dots, p_n$  can be regarded as guards allowing the monitor to nondeterministically transit to one of the states  $s_1, s_2, \dots, s_n$ .

**Definition 35** *Maps  $\theta: A \rightarrow \{\text{true}, \text{false}\}$  are called **A-events**, or simply **events**. Given an A-event  $\theta$ , we define its **multi-transition extension** as the map  $\theta_{MT}: MT(P_A, S) \rightarrow 2^S$ , where  $\theta_{MT}([p_1: s_1, p_2: s_2, \dots, p_n: s_n]) = \{s_i \mid \theta \models p_i\}$ .*

The role of A-events is to transmit the monitor information regarding the running program. In any program state, the map  $\theta$  assigns atomic propositions to *true* iff they hold in that state, otherwise to *false*. Therefore, A-events can be regarded as abstractions of the program states. Moreover, technically speaking, A-events are in a bijective map to  $P_A$ . For an *MT*  $\mu$ , the set of states  $\theta_{MT}(\mu)$  is often called the set of *possible continuations* of  $\mu$  under  $\theta$ .

**Example 14** *If  $\mu = [a \vee \neg b: s_1, \neg a \wedge b: s_2, c: s_3]$ , and  $\theta(a)=\text{true}$ ,  $\theta(b)=\text{false}$ , and  $\theta(c)=\text{true}$ , then the set of possible continuations of  $\mu$  under  $\theta$ ,  $\theta_{MT}(\mu)$ , is  $\{s_1, s_3\}$ .*



**Definition 36** A (*nondeterministic*) *binary transition tree (BTT)* over  $A$  and  $S$  is inductively defined as either a set in  $2^S$  or a structure of the form  $a ? \beta_1 : \beta_2$ , for some atom  $a$  and for some binary transition trees  $\beta_1$  and  $\beta_2$ . Let  $BTT(A, S)$  denote the set of *BTTs* over the set of states  $S$  and atoms  $A$ .

**Definition 37** Given an event  $\theta$ , we define its *binary transition tree extension* as the map  $\theta_{BTT} : BTT(A, S) \rightarrow 2^S$ , where:

$$\begin{aligned} \theta_{BTT}(Q) &= Q \text{ for any set of states } Q \subseteq S, \\ \theta_{BTT}(a ? \beta_1 : \beta_2) &= \theta_{BTT}(\beta_1) \text{ if } \theta(a) = \text{true}, \text{ and} \\ \theta_{BTT}(a ? \beta_1 : \beta_2) &= \theta_{BTT}(\beta_2) \text{ if } \theta(a) = \text{false}. \end{aligned}$$

**Definition 38** A *BTT*  $\beta$  *implements* an *MT*  $\mu$ , written  $\beta \models \mu$ , iff for any event  $\theta$ , it is the case that  $\theta_{BTT}(\beta) = \theta_{MT}(\mu)$ .

**Example 15** The *BTT*  $b ? (a ? (c ? s_1 s_3 : s_1) : (c ? s_2 : \emptyset)) : (c ? s_1 s_3 : s_3)$  implements the multi-transition shown in Example 14.

Fig. 9.2 represents this *BTT* graphically. The right branch of the node labeled with **b** corresponds to the *BTT* expression  $(c ? s_1 s_3 : s_3)$ , and similarly for the left branch and every other node. Atomic predicates can be any host programming language boolean expressions. For example, one may be interested if a variable  $x$  is positive or if a vector  $v[1...100]$  is sorted. Some atomic predicates typically are more expensive to evaluate than others. Since our purpose is to generate *efficient monitors*, we need to take the evaluation costs of atomic predicates into consideration. Moreover, some predicates can hold with higher probability than others; for example, some predicates may be simple “sanity checks”, such as checking whether the output of a

sorting procedure is indeed sorted. We next assume that atomic predicates are given evaluation costs and probabilities to hold. These may be estimated *apriori*, either statically or dynamically.

**Definition 39** *If  $\varsigma: A \rightarrow \mathcal{R}^+$  and  $\pi: A \rightarrow [0, 1]$  are cost and probability functions for events in  $A$ , respectively, then let  $\gamma_{\varsigma, \pi}: BTT(A, S) \rightarrow \mathcal{R}^+$  defined as:*

$$\begin{aligned} \gamma_{\varsigma, \pi}(Q) &= 0 \text{ for any } Q \subseteq S, \text{ and} \\ \gamma_{\varsigma, \pi}(a \text{ ? } \beta_1 : \beta_2) &= \varsigma(a) + \pi(a) * \gamma_{\varsigma, \pi}(\beta_1) + (1 - \pi(a)) * \gamma_{\varsigma, \pi}(\beta_2), \end{aligned}$$

be the **expected (evaluation) cost** function on *BTTs* in  $BTT(A, S)$ .

**Example 16** *Given  $\varsigma = \{(a, 10), (b, 5), (c, 20)\}$  and  $\pi = \{(a, 0.2), (b, 0.5), (c, 0.5)\}$ , the expected evaluation cost of the *BTT* defined in Example 2 is 30.*

With the terminology and motivations above, the following problem develops as an interesting and important problem in monitor synthesis:

**Problem:** Optimal  $BTT(A, S)$ .

**Input:** A multi-transition  $\mu = [p_1 : s_1, p_2 : s_2, \dots, p_n : s_n]$  with associated cost  $\varsigma : A \rightarrow \mathcal{R}^+$  and probability  $\pi : A \rightarrow [0, 1]$ .

**Output:** A minimal cost *BTT*  $\beta$  with  $\beta \models \mu$ .

Binary decision trees (BDTs) and diagrams (BDDs) have been studied as models and data-structures for several problems in artificial intelligence [95] and program verification [21]. Appendix B discusses BDTs and how they relate to *BTTs*. Moret [95] shows that a simpler version of this problem, using BDTs, is NP-hard.

In spite of this result, in general the number of atoms in formulae is relatively small, so it is not impractical to exhaustively search for the optimal *BTT*. We next informally describe a backtracking algorithm that we are currently using in our implementation to compute the minimal cost *BTT* by exhaustive search. Start with the sequence of all atoms in  $A$ . Pick one atom, say  $a$ , and make two recursive calls to this procedure, each assuming one boolean assignment to  $a$ . In each call, pass the remaining sequence of atoms to test, and simplify the set of propositions in the multi-transition according to the value of  $a$ . The product of the *BTTs* is taken when the recursive calls return in order to compute all *BTTs* starting with  $a$ . This procedure repeats until no atom is left in the sequence. We select the minimal cost *BTT* amongst all computed.

## 9.4 Binary Transition Tree Finite State Machines

We next define an automata-like structure, formalizing the desired concept of an *effective runtime monitor*. The transitions of each state are all-together encoded by a *BTT*, in practice the statistically optimal one, in order for the monitor to efficiently transit as events take place in the monitored program. Violations occur when one cannot further transit to any state for any event. A special state, called *neverViolate*, will denote a configuration in which one can no longer detect a violation, so one can stop the monitoring session if this state is reached.

**Definition 40** A *binary transition tree finite state machine (BTT-FSM)* is a tuple  $\langle A, S, btt, S_0 \rangle$ , where  $A$  is a set of atoms,  $S$  is a set of states potentially including a special state called “*neverViolate*”,  $btt$  is a map associating a *BTT* in  $BTT(A, S)$  to each state in  $S$  where  $btt(\text{neverViolate}) = \{\text{neverViolate}\}$  when  $\text{neverViolate} \in S$ , and  $S_0 \subseteq S$  is a subset of initial states.

**Definition 41** Let  $\langle A, S, btt, S_0 \rangle$  be a *BTT-FSM*. For an event  $\theta$  and  $Q, Q' \subseteq S$ , we write  $Q \xrightarrow{\theta} Q'$  and call it a **transition between sets of states**, whenever  $Q' = \bigcup_{s \in Q} \theta_{BTT}(btt(s))$ . A **trace of events**  $\theta_1 \theta_2 \dots \theta_j$  generates a **sequence of transitions**  $Q_0 \xrightarrow{\theta_1} Q_1 \xrightarrow{\theta_2} \dots \xrightarrow{\theta_j} Q_j$  in the *BTT-FSM*, where  $Q_0 = S_0$  and  $Q_i \xrightarrow{\theta_{i+1}} Q_{i+1}$ , for all  $0 \leq i < j$ . The trace is **rejecting** iff  $Q_j = \{\}$ .

Note that no finite extension of a trace  $\theta_1 \theta_2 \dots \theta_j$  will be **rejected** if  $\text{neverViolate} \in Q_j$ . The state *neverViolate* denotes a configuration in which violations can no longer be detected for any finite trace extension. This means that the set  $Q_k$  will not be empty, for any  $k > j$ , when  $\text{neverViolate} \in Q_j$ . Therefore, the monitoring session can stop at event  $j$  if  $\text{neverViolate} \in Q_j$ , because we are only interested in violations of requirements.

## 9.5 Generating a *BTT-FSM* from a Büchi automaton

Not any property can be monitored. For example, in order to check a liveness property one needs to ensure that certain propositions hold infinitely often, which cannot be verified at runtime. This section describes how to transform a Büchi automaton into an efficient *BTT-FSM* that rejects precisely the minimal bad prefixes of the denoted  $\omega$ -language.

**Definition 42** A **monitor FSM (MFSM)** is a tuple  $\langle \Sigma, S, \delta, S_0 \rangle$ , where  $\Sigma = P_A$  is an alphabet,  $S$  is a set of states potentially including a special state “*neverViolate*”,  $\delta: S \times \Sigma \rightarrow 2^S$  is a transition function with  $\delta(\text{neverViolate}, \text{true}) = \{\text{neverViolate}\}$  when  $\text{neverViolate} \in S$ , and  $S_0 \subseteq S$  are initial states.

Note that we take  $\Sigma$  to be  $P_A$ , the set of propositions over atoms in  $A$ . Like *BTT-FSMs*, *MFSMs* may also have a special *neverViolate* state.

**Definition 43** Let  $Q_0 \xrightarrow{\theta_1} Q_1 \xrightarrow{\theta_2} \dots \xrightarrow{\theta_j} Q_j$  be a sequence of transitions in the *MFSM*  $\langle \Sigma, S, \delta, S_0 \rangle$ , generated from  $t = \theta_1 \theta_2 \dots \theta_j$ , where  $Q_0 = S_0$  and  $Q_{i+1} = \bigcup_{s \in Q_i} \{\delta(s, \sigma) \mid \theta_{i+1} \models \sigma\}$ , for all  $0 \leq i < j$ . We say that the *MFSM* **rejects**  $t$  iff  $Q_j = \{\}$ .

No finite extension of  $t$  will be **rejected** if  $\text{neverViolate} \in Q_j$ .

**From Büchi to MFSM.** We next describe two simplification procedures on a Büchi automaton that are sound w.r.t. monitoring, followed by the construction of an *MFSM*. The first procedure identifies segments of the automaton which cannot lead to acceptance and can therefore be safely removed. As we will show shortly, this step is necessary in order to guarantee the soundness of the monitoring procedure. The second simplification identifies states with the property that if they are reached then the corresponding requirement cannot be violated by any *finite* extension of the trace, so monitoring is ineffective from there on. Note that reaching such a state does not necessarily mean that a good prefix has been recognized, but only that the property is *not monitorable* from there on.

**Definition 44** Let  $\langle \Sigma, S, \delta, S_0, \mathcal{F} \rangle$  be a Büchi automaton,  $C$  a connected component of its associated graph, and  $\text{nodes}(C)$  the states associated to  $C$ . We say that  $C$  is **isolated** iff for any  $s \in \text{nodes}(C)$  and  $\sigma \in \Sigma$ , it is the case that  $\delta(s, \sigma) \subseteq \text{nodes}(C)$ . We say that  $C$  is **total** iff for any  $s \in \text{nodes}(C)$  and event  $\theta$ , there are transitions  $\sigma$  such that  $\theta \models \sigma$  and  $\delta(s, \sigma) \cap \text{nodes}(C) \neq \emptyset$ .

Therefore, there is no way to escape from an isolated connected component, and regardless of the upcoming event, it is always possible to transit from any node of a total connected component to another node in that component.

**Removing Bad States.** The next procedure removes states of the Büchi automaton which cannot be part of any accepting run (see Definition 2).

Note that any state appearing in such an accepting run must eventually *reach* an accepting state. This procedure is fundamentally inspired by strongly-connected-component-analysis [75, 129], used to check emptiness of the language denoted by a Büchi automaton. Given a Büchi automaton  $\mathcal{A} = \langle \Sigma, S, \delta, S_0, \mathcal{F} \rangle$ , let  $U \subseteq S$  be the largest set of states such that the language of  $\langle \Sigma, S, \delta, U, \mathcal{F} \rangle$  is empty. The states in  $U$  are unnecessary in  $\mathcal{A}$ , because they cannot change its language. Fortunately,  $U$  can be calculated effectively as the set of states that *cannot reach* any cycle in the graph associated to  $\mathcal{A}$  which contains at least one accepting state in  $\mathcal{F}$ . Fig. 9.3 shows an algorithm to do this.

<p>INPUT : A Büchi automaton <math>\mathcal{A}</math>  OUTPUT : A smaller Büchi automaton <math>\mathcal{A}'</math> such that <math>\mathcal{L}(\mathcal{A}') = \mathcal{L}(\mathcal{A})</math>.  REMOVE_BAD_STATES :      for each maximal connected component <math>C</math> of <math>\mathcal{A}</math>          if (<math>C</math> is isolated and <math>nodes(C) \cap \mathcal{F} = \emptyset</math>) then mark all states in <math>C</math> “bad”  DFS_MARK_BAD ; REMOVE_BAD</p>
--

Figure 9.3: Removing bad states

The loop identifies maximal isolated connected components which do not contain any accepting states. The nodes in these components are marked as “bad”. The procedure DFS\_MARK\_BAD performs a depth-first-search in the graph and marks nodes as “bad” when all outgoing edges lead to a “bad” node. Finally, the procedure REMOVE\_BAD removes all the bad states. The runtime complexity of this algorithm is dominated by the computation of maximal connected components. In our implementation, we used Tarjan’s  $O(V + E)$  double DFS [21]. The proof of correctness is simple and it appears in Appendix A. The Büchi automaton  $\mathcal{A}'$  produced by the algorithm in Fig. 9.3 has the property that there is some proper path from any of its states to some accepting state. One can readily generate an *MFSM* from a Büchi automaton  $\mathcal{A}$  by first applying the procedure REMOVE\_BAD\_STATES in Fig. 9.3, and then ignoring the acceptance conditions.

**Theorem 18** *The MFSM generated from a Büchi automaton  $\mathcal{A}$  as above rejects precisely the minimal bad prefixes of  $\mathcal{L}(\mathcal{A})$ .*

**Proof:** Let  $\mathcal{A} = \langle \Sigma, S, \delta, S_0, \mathcal{F} \rangle$  be the original Büchi automaton, let  $\mathcal{A}' = \langle \Sigma, S', \delta', S'_0, \mathcal{F} \rangle$  be the Büchi automaton obtained from  $\mathcal{A}$  by applying the

algorithm in Fig. 9.3, and let  $\langle \Sigma, S', \delta', S'_0 \rangle$  be the corresponding *MFSM* of  $\mathcal{A}'$ . For any finite trace  $t = \theta_1 \dots \theta_j$ , let us consider its corresponding sequence of transitions in the *MFSM*  $Q_0 \xrightarrow{\theta_1} \dots \xrightarrow{\theta_j} Q_j$ , where  $Q_0$  is  $S'_0$ . Note that the trace  $t$  can also be regarded as a sequence of letters in the alphabet  $\Sigma$  of  $\mathcal{A}$ , because we assumed  $\Sigma$  is  $P_A$  and because there is a bijection between propositions in  $P_A$  and  $A$ -events. All we need to show is that  $t$  is a bad prefix of  $\mathcal{A}'$  if and only if  $Q_j = \emptyset$ . Recall that  $\mathcal{A}'$  has the property that there is some non-empty path from any of its states to some accepting state. Thus, one can build an infinite path in  $\mathcal{A}'$  starting with any of its nodes, with the property that some accepting state occurs infinitely often. In other words,  $Q_j$  is not empty iff the finite trace  $t$  is the prefix of some infinite trace in  $\mathcal{L}(\mathcal{A}')$ . This is equivalent to saying that  $Q_j$  is empty iff the trace  $t$  is a bad prefix in  $\mathcal{A}'$ . Since  $Q_j$  empty implies  $Q_{j'}$  empty for any  $j > j'$ , it follows that the *MFSM* rejects precisely the minimal bad prefixes of  $\mathcal{A}$ .  $\square$

Theorem 1 says that the *MFSM* obtained from a Büchi automaton as above can be used as a monitor for the corresponding  $\omega$ -language. Indeed, one only needs to maintain a current set of states  $Q$ , initially  $S'_0$ , and transform it accordingly as new events  $\theta$  are generated by the observed program: if  $Q \xrightarrow{\theta} Q'$  then set  $Q$  to  $Q'$ ; if  $Q$  ever becomes empty then report violation. Theorem 1 tells us that a violation will be reported as soon as a bad prefix is encountered.

**Collapsing Never-Violate States.** Reducing runtime overhead is crucial in runtime verification. There are many situations when the monitoring process can be safely stopped, because the observed finite trace cannot be finitely extended to any bad prefix. The following procedure identifies states in a Büchi automaton which cannot lead to the violation of any *finite* computation. For instance, the Büchi automaton in Fig. 9.4 can only reject infinite words in which the state  $s_2$  occurs finitely many times; moreover, at least one transition is possible at any moment. Therefore, the associated *MFSM* will never report a violation, even though there are infinite words that are not accepted. We call such an automaton *non-monitorable*. This example makes it clear that if a state like  $s_1$  is ever reached by the monitor, it does *not* mean that we found a good prefix, but that we could *stop looking* for bad prefixes.

Let  $\mathcal{A} = \langle \Sigma, S, \delta, S_0, \mathcal{F} \rangle$  be a Büchi automaton *simplified with REMOVE\_BAD\_STATES*. The procedure in Fig. 9.5 finds states which, if reached by a monitor, then the monitor can no longer detect violations regardless of what events will be observed in the future. The procedure first identifies the total connected

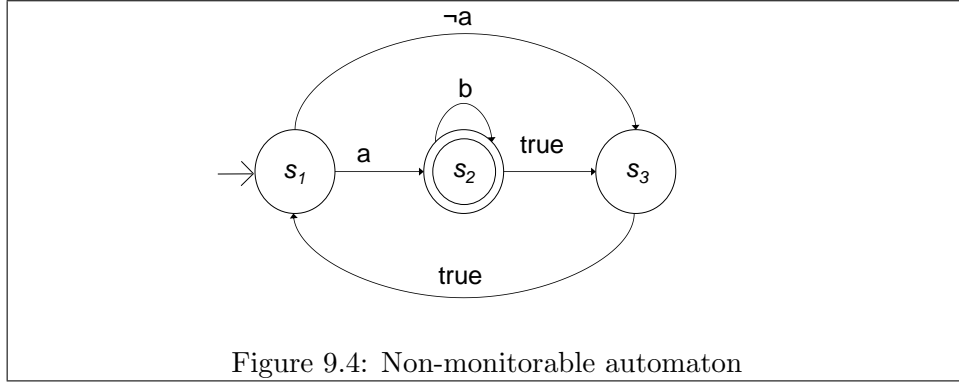


Figure 9.4: Non-monitorable automaton

components. According to the definition of totality, once a monitor reaches a state of a total connected component, the monitor will have the possibility to always transit within that connected component, thus never getting a chance to report violation. All states of a total component can therefore be marked as “never violate”. Other states can also be marked as such if, for any events, it is possible to transit from them to states already marked “never violate”; that is the reason for the disjunction in the second conditional. The procedure finds such nodes in a depth-first-search. Finally, `COLLAPSE-NEVER-VIOLATE` collapses all components marked “never violate”, if any, to a distinguished node, *neverViolate*, having just a *true* transition to itself. If any collapsed node was in the initial set of states, then the entire automaton is collapsed to *neverViolate*. The procedure `GENERATE_MFSM` produces an *MFSM* by ignoring accepting conditions.

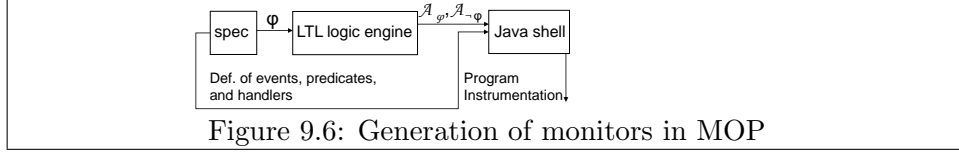
```

INPUT  : A Büchi automaton  $\mathcal{A}$ , cost function  $\varsigma$ , and probability function  $\pi$ .
OUTPUT : An effective BTT-FSM monitor rejecting the bad prefixes of  $\mathcal{L}(\mathcal{A})$ .
COLLAPSE_NEVER_VIOLATE :
  for each maximal connected component  $C$  of  $\mathcal{A}$ 
    if (  $C$  is total ) then mark all states in  $C$  as “never violate”
  for each  $s$  in depth-first-search visit
    if (  $\bigvee \{ \sigma \mid \delta(s, \sigma) \text{ contains some state marked “never violate”} \}$  )
      then mark  $s$  as “never violate”
  COLLAPSE_NEVER_VIOLATE ; GENERATE_MFSM ; GENERATE_BTT_FSM

```

Figure 9.5: Collapsing non-monitorable states

Taking as input this *MFSM*, say  $\langle \Sigma, S', \delta', S'_0 \rangle$ , cost function  $\varsigma$ , and probability function  $\pi$ , `GENERATE_BTT-FSM` constructs a *BTT-FSM*  $\langle A, S', btt, S'_0 \rangle$ , where  $A$  corresponds to the set of atoms from which the alphabet  $\Sigma$  is built, and the map *btt*, here represented by a set of pairs, is defined as follows:



$$\begin{aligned}
btt = & \{(neverViolate, \{neverViolate\}) \mid neverViolate \in S'\} \cup \\
& \{(s, \beta_s) \mid s \in S' - \{neverViolate\} \wedge \beta_s \models \mu_s\}, \text{ where} \\
& \beta_s \text{ optimally implements } \mu_s \text{ w.r.t. } \varsigma \text{ and } \pi, \text{ with } \mu_s = \oplus(\bigcup\{[\sigma : s'] \mid s' \in \delta'(s, \sigma)\})
\end{aligned}$$

The symbol  $\oplus$  denotes concatenation on a set of multi-transitions. Optimal *BTTs*  $\beta_s$  are generated like in Section 9.3. Proof of correctness appears in Appendix A.

## 9.6 Monitor Generation and MOP

We have shown that one can generate from a Büchi automaton a *BTT-FSM* recognizing precisely its bad prefixes. However, it is still necessary to integrate the *BTT-FSM* monitor within the program to be observed. *Monitoring-oriented programming* (MOP) [18] aims at merging specification and implementation through generation of runtime monitors from specifications and integration of those within implementation. In MOP, the task of generating monitors is divided into defining a logic engine and a language shell. The logic engine is concerned with the translation of specifications given as logical formulae into monitoring (pseudo-)code. The shell is responsible for the integration of the monitor within the application.

Fig. 9.6 captures the essence of the synthesis process of LTL monitors in MOP using the technique described in this paper. The user defines specifications either as annotations in the code or in a separate file. The specification contains definitions of events and state predicates, as well as LTL formulae expressing trace requirements. These formulae treat events and predicates as atomic propositions. Handlers are defined to track violation or validation of requirements. For instance, assume the events  $a$  and  $b$  denote the login and the logoff of the same user, respectively. Then the formula  $\Box(a \rightarrow \circ(\neg a \mathcal{U} b))$  states that the user cannot be logged in more than once. A violation handler could be declared to track the user who logged in twice. The logic engine is responsible for the translation of the formulae  $\varphi$  and  $\neg\varphi$  into two *BTT-FSM* monitors. One detects violation and the other validation of  $\varphi$ . Note that if the user is just interested in validation (no violation handler), then only the automaton for negation is generated. Finally, the

Temporal Formula	<i>BTT-FSM</i>
$\Box(a \rightarrow b \mathcal{U} c)$	$[s_0] \rightsquigarrow c ? (b ? s_0 s_1 : s_0) : (a ? (b ? s_1 : \emptyset) : (b ? s_0 s_1 : s_0))$ $s_1 \rightsquigarrow b ? (c ? s_0 s_1 : s_1) : (c ? s_0 : \emptyset)$
$\Box(a \rightarrow \Diamond b)$	$[neverViolate] \rightsquigarrow \{neverViolate\}$
$a \mathcal{U} b \mathcal{U} c$	$[s_0] \rightsquigarrow c ? neverViolate : (a ? (b ? s_0 s_1 : s_0) : (b ? s_1 : \emptyset))$ $s_1 \rightsquigarrow c ? neverViolate : (b ? s_1 : \emptyset)$ $neverViolate \rightsquigarrow \{neverViolate\}$

Table 9.1: *BTT-FSMs* generated from temporal formulae

language shell reads the definition of events and instruments the code so that the monitor will receive the expected notifications.

We used LTL2BA [97] to generate standard Büchi automata from LTL formulae. The described procedures are implemented in JAVA. This software and a WWW demo are available from the MOP website [1].

### 9.6.1 Evaluation

Table 1 shows *BTT-FSM* monitors for some LTL formulae. The *BTT* definition corresponding to a state follows the arrow ( $\rightsquigarrow$ ). Initial states appear in brackets. For producing this table, we used the same cost and probabilities for all events and selected the smallest *BTT*. The first formula cannot be validated by monitoring and presents the permanent possibility to be violated; that is why its *BTT-FSM* does not have a *neverViolate* state. The second formula can never be violated since event *a* followed by event *b* can always occur in the future, so its *BTT-FSM* consists of just one state *neverViolate*. The last formula shows that our procedure does not aim at distinguishing validating from non-violating prefixes.

Table 2 shows that our technique can not only identify non-monitorable formulae, but also reduce the cost of monitoring by collapsing large parts of the Büchi automaton. We use the symbols  $\heartsuit$ ,  $\clubsuit$ , and  $\spadesuit$  to denote, respectively, the effectiveness of REMOVE\_BAD\_STATES, the first, and the second loop of COLLAPSE\_NEVER\_VIOLATE. The first group contains non-monitorable formulae. The next contains formulae where monitor size could not be reduced by our procedures. The third group shows formulae where our simplifications could significantly reduce the monitor size. The last group shows examples of “accidentally” safe and “pathologically” safe formulae from [83]. A formula  $\varphi$  is accidentally safe iff not all bad prefixes are “informative”












Temporal Formula	# states	# transitions	symlif.
$\diamond a$	2 , 1	3 , 1	
$a \mathcal{U} \circ(\diamond b)$	3 , 1	5 , 1	
$\square(a \wedge b \rightarrow \diamond c)$	2 , 1	4 , 1	
$a \mathcal{U} (b \mathcal{U} (c \mathcal{U} (\diamond d)))$	2 , 1	3 , 1	
$a \mathcal{U} (b \mathcal{U} (c \mathcal{U} \square(d \rightarrow \diamond e)))$	5 , 1	15 , 1	
$\neg a \mathcal{U} (b \mathcal{U} (c \mathcal{U} \square(d \rightarrow \diamond e)))$	12 , 1	51 , 1	
$\neg \diamond a$	1 , 1	1 , 1	
$\square(a \rightarrow b \mathcal{U} c)$	2 , 2	4 , 4	
$a \mathcal{U} (b \mathcal{U} (c \mathcal{U} d))$	4 , 4	10 , 10	
$a \wedge \circ(\diamond b) \wedge \diamond(\square(e))$	5 , 4	11 , 6	
$a \wedge \circ(\diamond b) \wedge \circ(\diamond c) \wedge \diamond(\square(e))$	9 , 6	29 , 12	
$a \wedge \circ(\diamond b) \wedge \circ(\diamond c) \wedge \circ(\diamond d) \wedge \diamond(\square(e))$	17 , 10	83 , 30	
$a \wedge \circ(\neg(\square(b \rightarrow c \mathcal{U} d))) \wedge \diamond(\square(e))$	7 , 5	20 , 10	
$\square(a \vee \circ(\square(c))) \wedge \square(b \vee \circ(\square(\neg c)))$	3 , 3	5 , 5	
$(\square(a \vee \diamond(\square(c))) \wedge \square(b \vee \diamond(\square(\neg c)))) \vee \square(a) \vee \square(b)$	12 , 6	43 , 22	

Table 9.2: Number of states and transitions before and after monitoring simplifications

[83] (i.e., can serve as a witness for violation) but all computations that violate  $\varphi$  have an informative bad prefix. A formula  $\varphi$  is pathologically safe if there is a computation that violates  $\varphi$  and has no informative bad prefix. Since we detect *all* minimal bad prefixes, informativeness does not play any role in our approach. Both formulae are monitorable. For the last formula, in particular, a minimal bad prefix will be detected as soon as the monitor observes a  $\neg a$ , having previously observed a  $\neg b$ . One can generate and visualize the *BTT-FSMs* of all these formulae, and many others, online at [1].

## 9.7 Conclusions

Not all properties a Büchi automaton can express are monitorable. This paper describes transformations that can be applied to extract the monitorable components of Büchi automata, reducing their size and the cost of runtime verification. The resulting automata are called *monitor finite state machines (MFSMs)*. The presented algorithms have polynomial running time in the size

of the original Büchi automata and have already been implemented. Another contribution of this paper is the definition and use of *binary transition trees* (*BTTs*) and corresponding finite state machines (*BTT-FSMs*), as well as a translation from *MFSMs* to *BTT-FSMs*. These special-purpose state machines encode optimal evaluation paths of boolean propositions in transitions.

We used LTL2BA [97] to generate Büchi automata from LTL, and JAVA to implement the presented algorithms. Our algorithms, as well as a graphical HTML interface, are available at [1]. This work is motivated by, and is part of, a larger project aiming at promoting monitoring as a foundational principle in software development, called *monitoring-oriented programming* (MOP). In MOP, the user specifies formulae, atoms, cost and probabilities associated to atoms, as well as violation and validation handlers. Then all these are used to automatically generate monitors and integrate them within the application.

This work is concerned with monitoring *violations* of requirements. In the particular case of LTL, *validations* of formulae can also be checked using the same technique by monitoring the negation of the input formula. Further work includes implementing the algorithm defined in [41] for generating Büchi automata of size  $2^{O(|\varphi|)}$  from PTL, combining multiple formulae in a single automaton as showed by Ezick [35] so as to reduce redundancy of proposition evaluations, and applying further (standard) NFA simplifications to *MFSM*.

## Appendix A. Proof of Correctness

(*Correctness of REMOVE\_BAD\_STATES*) Let  $\mathcal{A} = \langle \Sigma, S, \delta, S_0, \mathcal{F} \rangle$  be the input and  $\mathcal{A}' = \langle \Sigma, S', \delta', S'_0, \mathcal{F} \rangle$  the output Büchi automata of the procedure. We want to show that their denoted languages are the same. Let  $\rho$  be an accepting run of  $\mathcal{A}$ . Then  $\rho$  can be fragmented in  $\rho' \rho''$  where  $\rho'$  is the prefix whose states appear only finitely many times in  $\rho$ . Each state in  $\rho''$  is therefore reachable from any other and therefore must be in a connected component where some states are in  $\mathcal{F}$ . The converse is also true, i.e., any connected component reachable from the set of initial states having at least one state in  $\mathcal{F}$  generates accepting runs (from [21] pp. 129). It is then immediate to notice that no accepting run  $\rho$  contains states that can never reach a state in  $\mathcal{F}$ . It turns out that any state belonging to an isolated component with no accepting states can be trivially removed as well as any state that can *only* make transitions to others which never reach an accepting state. This reachability analysis is performed in a depth-first-search as usual. Since only states that will never appear in accepting runs of  $\mathcal{A}$  are removed, it follows that  $\mathcal{L}(\mathcal{A}) = \mathcal{L}(\mathcal{A}')$ .

(*Correctness of COLLAPSE\_NEVER\_VIOLATE*) First, we need to show that the *MFSM* produced by `GENERATE_MFSM` still detects bad prefixes of  $\mathcal{L}(\mathcal{A})$ , where  $\mathcal{A}$  is the input Büchi automaton having the property that all states can reach accepting states. From the observation that all states that have been collapsed to *neverViolate* can reach an accepting state, the proof of the first part is similar to that of Theorem 1 and is omitted. So, the *MFSM* is still a monitor for the bad prefixes of  $\mathcal{L}(\mathcal{A})$ . In other words, this simplification is sound w.r.t. monitoring minimal bad prefixes. In addition to this, we need to show that the state *neverViolate* is reached *only if* violations of the requirement can no longer occur. We prove this second correctness criteria by a case analysis on the shape of the *MFSM* associated to the input Büchi. Each case corresponds to a loop in `COLLAPSE_NEVER_VIOLATE`.

(case 1) Let  $Q_0 \xrightarrow{\theta_1} Q_1 \xrightarrow{\theta_2} \dots \xrightarrow{\theta_i} Q_j$  be a sequence of transitions on the trace  $\theta_1 \theta_2 \dots \theta_j$  produced by the *MFSM* corresponding to the input Büchi automaton  $\mathcal{A}$ . Recall that we assume `REMOVE_BAD_STATES` is called before `COLLAPSE_NEVER_VIOLATE`. That is, all states in the input automaton reach some state in  $\mathcal{F}$ . If some node in a total connected component  $C$  of the

graph associated to  $\mathcal{A}$  belongs to  $Q_j$  then it is not possible for any finite trace with prefix  $\theta_1\theta_2...\theta_j$  to be rejected by the corresponding *MFSM* since it is always possible to make a transition between nodes in  $C$  from the definition of totality. Those states in  $C$  can be marked as “never violate” meaning that they denote the special *neverViolate* state.

(case 2) Let  $Q_0 \xrightarrow{\theta_1} Q_1 \xrightarrow{\theta_2} \dots \xrightarrow{\theta_j} Q_j$  be a sequence of transitions produced on the trace  $\theta_1\theta_2...\theta_j$  by the *MFSM* defined as before. If any state belongs to  $Q_j$  with the property that for further events  $\theta_{j+1}$  there exists a transition to a state marked “never violate”, then there must exist some state marked “never violate” in  $Q_{j+1}$ . Such states can be found by checking the validity of the disjunction of propositions labeling the edges of transitions to states marked “never violate”. Since the monitor will definitely reach a state marked “never violate” in the trace  $\theta_1\theta_2...\theta_j\theta_{j+1}$ , it is therefore safe to reach “never violate” in  $Q_j$  as well, since violations are impossible from event  $j$  on.

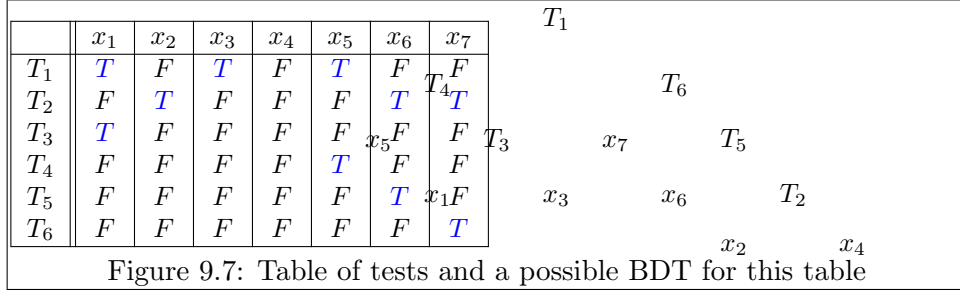
The states marked with the “never violate” label can therefore be collapsed, since a violation cannot occur if any of these states is reached. The *BTT-FSM* is generated from the *MFSM* according to the defined construction.

## Appendix B. Complexity Results

We use the definition of Binary Decision Trees (BDTs) due to Garey [40]. As we will show next, Garey’s BDTs serve as a procedure to identify one among a set of possible objects characterizing some aspect of interest. Hyafil and Rivest [72] proved the Optimal BDT problem NP-complete. Moret [95] shows that these BDTs are a restricted form of decision trees similar to binary transition trees as defined here.

Let  $X = \{x_1, \dots, x_n\}$  be a finite set of objects and  $\mathcal{T} = \{T_1, \dots, T_t\}$  a finite set of tests. For each test  $T_i$ ,  $1 \leq i \leq t$ , and object  $x_j$ ,  $1 \leq j \leq n$ , we either have  $T_i(x_j) = \text{true}$  or  $T_i(x_j) = \text{false}$ . A binary decision tree associates tests in the root and all other internal nodes, and associates objects in  $X$  to terminal nodes. The *optimal decision tree* problem is to determine whether there exists a decision tree with cost less than or equal to  $w$  which completely identifies each element in  $X$ , given  $\mathcal{T}$  and  $X$ . The cost of this tree is defined as  $\sum_{x_i \in X} p(x_i)$ , where  $p(x_i)$  is the length of the path from the root to the terminal identifying  $x_i$ .

One might think of the objects in  $X$  as possible answers to a question. The tests in  $\mathcal{T}$  serve to prune the data set of answers. A decision tree defines



possible sequences of questions to ask in order to give a unique answer among those in  $X$ . The table of Fig. 9.7 appears in [40] and denotes the set of tests available to use in some binary decision tree.

The definition above [40] and the NP-completeness proof for the Optimal BDT problem [72] assume the table provided as input is unambiguous and completely identifies the objects in  $X$ . That is, any object in  $X$  can be distinguished by applying some sequence of tests. The decision tree in Fig. 9.7 appears in [40] as an identification procedure for the tests and objects defined in the adjacent table. The subtree to the left of a node denotes objects whose answers to the test labeling that node are *true* and the result of all other tests applied in the path to the root are consistent. The right subtree is similar. This is similar to binary transition trees. However, it is worth noting that in order to identify  $x_j$  one does not need to test all  $T_i$  having  $T_i(x_j) = \text{true}$ . Observe this in particular for  $x_6$  and  $x_7$ . It means that decision trees prune the data set in each test they make. A decision can be made whenever it is possible to identify an object. For instance, no object but  $x_7$  assigns *false* to  $T_1$  and *true* to  $T_6$ . So this is a sufficient condition to identify  $x_7$ . Moret showed that the construction of optimal binary decision trees, similar to *BTTs* (where tests take the form of boolean proposition over a set of atoms) is an NP-hard [95] problem. That implies that our *BTT* problem is also NP-hard.

## Chapter 10

# More Stuff

### 10.1 Optimal Monitoring of “Always Past” Temporal Safety

A monitor synthesis algorithm from linear temporal logic (LTL) safety formulae of the form  $\Box\varphi$  where  $\varphi$  is a past time LTL formula was presented in [59]. The generated monitors implemented the recursive semantics of past-time LTL using a dynamic programming technique, and needed  $O(|\varphi|)$  time to process each new event and  $O(|\varphi|)$  total space. Some compiler-like optimizations of the generated monitors were also proposed in [59], which would further reduce the required space. It is not clear how much the required space could be reduced by applying those optimizations.

We here show how to generate using a divide-and-conquer technique directly monitors that need  $O(k)$  space and still  $O(|\varphi|)$  time, where  $k$  is the number of temporal operators in  $\varphi$ .

#### 10.1.1 The Monitor Synthesis Algorithm

For simplicity, we assume only two past operators, namely  $\odot$  (previously) and  $\mathcal{S}$  (since). Let us first note that one cannot asymptotically reduce the space requirements below  $\Omega(k)$ , where  $k$  is the number of temporal operators appearing in the formula to monitor  $\varphi$ . Indeed, one can take  $\varphi = (\#_1 \rightarrow t_1) \wedge \cdots \wedge (\#_k \rightarrow t_k)$ , where for each  $1 \leq i \leq k$ ,  $\#_i$  is some event and  $t_i$  is some temporal formula containing precisely one past temporal operator, i.e., a  $\odot$  or a  $\mathcal{S}$ . Any monitor for  $\varphi$  must directly or indirectly store the status of each  $t_i$  at every event, to be able to react accordingly in

case the next event is some  $\#_i$ . Assuming that the events  $\#_i$  are distinct and that the formulae  $t_i$  are unrelated, then the monitor needs to distinguish among  $2^k$  possible states, so it needs  $\Omega(k)$  space.

In what follows, we assume the usual recursive semantics of LTL, also presented below, restricted to safety formulae of the form  $\Box\varphi$ , where  $\varphi$  is a past-time LTL. We adopt the simplifying assumption that the empty trace invalidates any atomic proposition and any past temporal operator; as argued in [59], this may not always be the best choice, but other semantic variations regarding the empty trace present no difficulties for monitoring.

**Definition 45** (adapted from [89]) *LTL formulae of the form  $\Box\varphi$  (read “always  $\varphi$ ”), where  $\varphi$  is a past-time LTL formula, are called LTL safety formulae; we may call them just safety formulae when LTL is understood from the context. An infinite trace  $u \in \Sigma^\omega$  satisfies  $\Box\varphi$ , written  $u \models \Box\varphi$ , iff each  $w \in \text{prefixes}(u)$  satisfies the past-time LTL formula  $\varphi$ , written also  $w \models \varphi$  and defined inductively as follows:*

$w \models \text{true}$	<i>is always true,</i>
$ws \models a$	<i>iff <math>a(s)</math> holds,</i>
$w \models \neg\varphi$	<i>iff <math>w \not\models \varphi</math>,</i>
$w \models \varphi_1 \wedge \varphi_2$	<i>iff <math>w \models \varphi_1</math> and <math>w \models \varphi_2</math>,</i>
$ws \models \circ\varphi$	<i>iff <math>w \models \varphi</math>,</i>
$ws \models \varphi_1 \mathcal{S} \varphi_2$	<i>iff <math>ws \models \varphi_2</math> or <math>ws \models F\varphi</math> and <math>w \models \varphi_1 \mathcal{S} \varphi_2</math></i>
$\epsilon \models \varphi$	<i>is false otherwise</i>

Given safety formula  $\Box\varphi$ , we let  $\mathcal{L}(\Box\varphi) \subseteq \Sigma^\omega$  be the set  $\{u \in \Sigma^\omega \mid u \models \Box\varphi\}$ .

**Proposition 16**  $\mathcal{L}(\Box\varphi) \in \text{Safety}^\omega$  for any past-time LTL formula  $\varphi$ .

**Proof:** By the definition of  $\mathcal{L}(\Box\varphi)$  in Definition 45 and the definition of  $\Box P$  in Definition 12, one can easily note that  $\mathcal{L}(\Box\varphi) = \Box\mathcal{L}(\varphi)$ , where  $\mathcal{L}(\varphi) = \{w \in \Sigma^* \mid w \models \varphi\}$ . Therefore,  $\mathcal{L}(\Box\varphi) \in \text{Safety}_\Box^\omega$ . The rest follows by Theorem 5.  $\square$

Let us next investigate the problem of monitoring safety properties  $P \in \text{Safety}^\omega$  expressed as languages of safety formulae, that is,  $P = \mathcal{L}(\Box\varphi)$  for some past-time LTL formula  $\varphi$ . Because of the recursive nature of the satisfaction relation, a first important observation is that the generated monitor only needs to store information regarding the status of temporal operators from the previous state. More precisely, the monitor needs one bit per temporal operator, keeping the satisfaction status of the subformula

corresponding to that temporal operator; when a new state is received, the satisfaction status of the subformula is recalculated according to the recursive semantics above and then the bit is updated. The order in which the temporal operators are processed when a new state is received is important: the nested operators must be processed first.

We next present the actual monitor synthesis algorithm at a high-level. We refrain from giving detailed pseudocode as we did in [59], because different applications may choose different implementation paradigms. For example, we are currently using rewriting techniques to implement the monitor synthesis algorithms in MOP [20]; Section 10.1.2 shows our complete Maude rewriting implementation of the subsequent monitor synthesis algorithm.

**Step 1** Let  $\varphi_1, \dots, \varphi_k$  be the  $k$  subformulae of  $\varphi$  corresponding to temporal operators, such that, if  $\varphi_i$  is a subformula of  $\varphi_j$ , then  $i < j$ ; this can be easily achieved by a DFS traversal of  $\varphi$ .

**Step 2** Let  $bit[1..k]$  be a vector of  $k$  bits initialized to 0 (or false);  $bit[i]$  will store information related to  $\varphi_i$  from the previous state:

- if  $\varphi_i = \circ\psi$  then  $bit[i]$  says if  $\psi$  was satisfied at the previous state;
- if  $\varphi_i = \psi \mathcal{S} \psi'$  then  $bit[i]$  says if  $\varphi_i$  was satisfied at the previous step.

**Step 3** Let  $bit'[1..k]$  be another vector of  $k$  bits; this will be used to store temporary results, which will be moved eventually into the vector  $bit[1..k]$ .

**Step 4** Generate a loop that executes whenever a new state  $s$  is available; the body of the loop executes the following code:

**Step 4.1** For each  $i$  from 1 to  $k$  execute a bit assignment as follows, where for a subformula  $\psi$  of  $\varphi$ ,  $\overline{\psi}$  is the boolean expression replacing in  $\psi$  each non-nested temporal subformula  $\varphi_j$  by  $bit[j]$  if  $\varphi_j$  is a “previously” formula or by  $bit'[j]$  if  $\varphi_j$  is a “since” formula, and each remaining atomic proposition  $a$  by its satisfaction in the current state,  $a(s)$ :

- if  $\varphi_i = \circ\psi$  then generate the assignment  $bit'[i] := \overline{\psi}$
- if  $\varphi_i = \psi \mathcal{S} \psi'$  then generate the assignment  $bit'[i] := \overline{\psi'} \vee \overline{\psi} \wedge bit[i]$

**Step 4.2** Generate the conditional: if  $\bar{\varphi}$  is false then error (formula violated)

**Step 4.3** Generate code to move the contents of  $bit'[1..k]$  into  $bit[1..k]$ .

Note that the generated monitors are well-defined, because each time a  $\bar{\psi}$  boolean expression is generated, all the bits in  $bit'[1..k]$  that are needed are already calculated. One can also perform boolean simplifications when calculating  $\bar{\psi}$  to reduce runtime overhead even further. For example, in our implementation that also generated the code below (see Section 10.1.2), we used the simplification  $\neg\neg\psi = \psi$ . To illustrate the monitor generation algorithm above, let us consider the past time LTL formula:  $\varphi = \neg(a \wedge \neg(\odot b \wedge (c \mathcal{S} (d \wedge (\neg e \mathcal{S} f))))$ . Step 1 produces the following enumeration of  $\varphi$ 's subformulae:  $\varphi_1 = \odot b$ ,  $\varphi_2 = \neg e \mathcal{S} f$ , and  $\varphi_3 = c \mathcal{S} (d \wedge (\neg e \mathcal{S} f))$ . The other steps eventually generate the code:

```

bit[1..3] := false;      // three global bits
foreach new state s do {
    // first update the bits in a consistent order
    bit'[1] := b(s);
    bit'[2] := f(s)  $\vee$  ( $\neg e(s) \wedge bit[2]$ );
    bit'[3] := d(s)  $\wedge$  bit'[2]  $\vee$  (c(s)  $\wedge$  bit[3]);
    // then check whether the formula is violated
    if a(s)  $\wedge$   $\neg(bit[1] \wedge bit'[3])$  then Error;
    // finally, update the state of the monitor
    bit[1..3] := bit'[1..3]
}

```

It is easy to see that for any past LTL formula  $\varphi$  of  $k$  temporal operators, the state of the generated monitor is encoded on  $k$  bits, namely the vector  $bit[1..k]$ . The runtime of the generated monitor is still  $O(|\varphi|)$ , because each temporal operator in  $\varphi$  results in an assignment and a read operation in the monitor, while each boolean operator in  $\varphi$  is “executed” by the monitor.

### 10.1.2 A Maude Implementation of the Monitor Synthesizer

We here show a term rewriting implementation of the algorithm above, using the Maude system [25]. Implementations in other languages are obviously also possible; however, rewriting proved to be an elegant means to generate monitors from logical formulae in several other contexts, and so seems to be here. In what follows we show the complete Maude code that takes as input

a formula, parses it, generates the monitor, and then pretty prints it. We use the K technique here [107], which is a rewriting-based language and/or logic definitional technique; to use K, one needs to first upload the generic, i.e., application-independent, module discussed at the end of this section.

## Atomic Predicates

We start by defining the atomic state predicates that one can use in formulae. These can be either identifiers (of the form 'a, 'abc, 'a123, etc.; these are provided by the Maude builtin module QID):

```
fmod PREDICATE is
  --- atomic predicates can be quoted identifiers
  protecting QID .
  sort Predicate .
  subsort Qid < Predicate .
endfm
```

## Syntax of Formulae

Let us next define the syntax of formulae. We here use Maude's mixfix notation for defining syntax as algebraic operators, where underscores stay for arguments. Also, note that operators are assigned precedences (declared as operator attributes), to relive the user from writing parentheses (the lower the precedence the tighter the binding):

```
fmod SYNTAX is
  protecting PREDICATE .
  sort Formula .
  subsort Predicate < Formula .
  op !_ : Formula -> Formula [prec 20] .
  op _/\_ : Formula Formula -> Formula [prec 23] .
  op 0_ : Formula -> Formula [prec 21] .
  op _S_ : Formula Formula -> Formula [prec 22] .
endfm
```

## Target Language

We are done with the input language. Let us now define the output language. We need a very simple language for implementing the generated monitors, namely one with limited assignment, conditional and looping. The generated code, as well as the target language, play no role in this paper; one is expected to change the language below to one's desired target language (Java, C,

C#, assembler, etc.). Our chosen language below has bits, expressions, statements and code. Bits are also expressions; code is a list of statements composed sequentially using “;” or just concatenation. The syntax below is also making use of precedence attributes. The `format` attributes are useful solely for pretty-printing reasons (see Maude’s manual [25] for details on formatting):

```
fmod CODE is
  --- syntax for the generated code
  protecting PREDICATE + INT + STRING .
  sorts Bit Exp Statement Code .
  subsorts Bit < Exp .
  subsort Statement < Code .
  ops (bit[_]) (bit'[_]) : Nat -> Bit .
  ops (bit[1 .. _]) (bit'[1 .. _]) : Int -> Bit .
  op _(s) : Predicate -> Exp [prec 0] .
  ops true false : -> Exp .
  op !_ : Exp -> Exp [prec 20] .
  op _/\_ : Exp Exp -> Exp [prec 23] .
  op _\/_ : Exp Exp -> Exp [prec 24] .
  op _:=_ : Exp Exp -> Statement [prec 27 format(ni d d d)] .
  op if_then_ : Exp Statement -> Statement
    [format(ni d d ++ --) prec 30] .
  op foreach new state s do _ : Code -> Statement
    [format(n d d d d s++ --n)] .
  op Error : -> Statement [format(ni d)] .
  op //_ : String -> Statement [format(ni d d)] .
  op nil : -> Code .
  op _;_ : Code Code -> Code [assoc id: nil prec 40] .
  op __ : Code Code -> Code [assoc id: nil prec 40] .
  op {_} : Code -> Statement [format(d d --ni ++)] .
  --- code simplification rules
  var B : Exp .
  eq !! B = B .
endfm
```

The following module defines the actual monitor synthesis algorithm. We use the K definitional technique here, because it yields a very compact implementation. K is centered on the basic intuition of *computation*; computations are encoded as first-order data-structures that “evolve”, via rewriting, to *results*. Computations are sequentialized using the list constructor “`_->_`”; thus, if K and K’ are computations, then `K -> K’` is the computation consisting of K followed by K’. Computations may eventually yield results; for example, `K -> K’` may rewrite (in context) to `R -> K’`, meaning that R is the result that K reduces to. An important feature of K is that one can

schedule lists of tasks for reduction; for example,  $[K1, K2, K3] \rightarrow K$  may eventually reduce to  $[R1, R2, R3] \rightarrow K$ , where  $R1$ ,  $R2$ , and  $R3$  are the results that  $K1$ ,  $K2$ , and  $K3$  reduce to, in this order. To use  $K$ , one needs to import the module  $K$  discussed at the end of this section. The equations of the module  $K$  (three in total) are all about reducing a list of computations to a list of results, supposing that one knows how to reduce one computation to one result.

$K$  is a definitional framework that is generic in computations and results. More precisely, it provides sorts `KComputation` and `KResult`, and expects its user to define the desired computations and results, as well as rules to reduce a computation to a result. Computations typically can be reduced to results only in context; to facilitate this,  $K$  provides a sort `KConfiguration`, which is also supposed to be populated accordingly. The sort `KConfiguration` is a multi-set sort over a sort `KConfigurationItem`, where the multi-set constructor is just concatenation; also, the sort `KComputation` is a list sort over `KComputationItem`, where the list constructor is `_->_`. To make use of  $K$ , one needs to first define constructors for the sorts `KConfigurationItem`, `KComputationItem` and `KResult`, and then to define how each computation item reduces to a result.

In our case, the computations are the formulae or subformulae that still need to be processed, and the results are the corresponding boolean expressions that need to be checked in the current (generated code) context to see whether the formula has been violated or not. We define the following additional constructors: we add four constructors for configurations, namely “`k`” that wraps the current computation, “`code`” that wraps the current generated code, and “`nextBit`” that wraps the next available bit; we add one main constructor for computations, “`form`”, that wraps a formula, and one constant computation item per operator in the input language (the later is needed to know how to combine back the results of the corresponding subexpressions; finally, we add one constructor for results, “`exp`”, that wraps a boolean expression.

The formula is processed in a depth-first-order, following a divide-and-conquer philosophy. Each subformula is decomposed into a list of computation subtasks consisting of its subformulae, then the corresponding results are composed back into a result corresponding to the original subformula. Recall that equations/rules apply wherever they match, not only at the top. Let us only discuss the two equations defining the “since” (`_S_`), the last two in the module below. The first one is straightforward: it decomposes

the task of processing  $F1 \ S \ F2$  to the subtasks of processing  $F1$  and  $F2$ ; the computation item  $S$  is placed in the computation structure to prepare the terrain for the next equation. The next equation applies after  $F1$  and  $F2$  have been processed, say to expressions  $B1$  and  $B2$ , respectively; if  $C$  is the code generated so far and if  $I+1$  is the next bit available, then the boolean expression corresponding to the current since formula is indeed  $\text{bit}'(I+1)$ , provided that one adds the corresponding code capturing the recursive semantics of since to the generated code.

```
fmod MONITOR-GENERATION is
  protecting K + SYNTAX + CODE .
  op k : KComputation -> KConfigurationItem .
  op code : Code -> KConfigurationItem .
  op nextBit : Nat -> KConfigurationItem .
  op process : Formula -> KConfiguration .
  op form : Formula -> KComputationItem .
  op exp : Exp -> KResult .
  ops ! /\ 0 S : -> KComputationItem .
  var P : Predicate . vars F F1 F2 : Formula . var C : Code .
  var I : Nat . vars B B1 B2 : Exp . var K : KComputation .
  eq process(F) = k(form(F)) code(nil) nextBit(0) .
  eq k(form(P) -> K) = k(exp(P(s)) -> K) .
  eq form(! F) = form(F) -> ! .
  eq exp(B) -> ! = exp(! B) .
  eq form(F1 /\ F2) = [form(F1),form(F2)] -> /\ .
  eq [exp(B1),exp(B2)] -> /\ = exp(B1 /\ B2) .
  eq form(0 F) = form(F) -> 0 .
  eq k(exp(B) -> 0 -> K) code(C) nextBit(I)
    = k(exp(bit[I + 1]) -> K) code(C ; bit'[I + 1] := B)
      nextBit(I + 1) .
  eq form(F1 S F2) = [form(F1), form(F2)] -> S .
  eq k([exp(B1),exp(B2)] -> S -> K) code(C) nextBit(I)
    = k(exp(bit'[I + 1]) -> K)
      code(C ; bit'[I + 1] := B2 \/ B1 /\ bit[I + 1]) nextBit(I + 1) .
endfm
```

## Putting It All together

The following module plugs the code generated above into the general pattern:

```
fmod PRETTY-PRINT is
  protecting MONITOR-GENERATION .
  sort Monitor .
  op genMonitor : Formula -> Code .
```

```

op makeMonitor : KConfiguration -> Code .

var F : Formula . var B : Exp . var C : Code . vars N M : Nat .
eq genMonitor(F) = makeMonitor(process(F)) .
eq makeMonitor(k(exp(B)) code(C) nextBit(N))
  = bit[1 .. N] := false ;
  foreach new state s do {
    // "first update the bits in a consistent order"
    C ;
    // "then check whether the formula is violated"
    if !(B) then Error ;
    // "finally, update the state of the monitor"
    bit[1 .. N] := bit'[1 .. N]
  } .
endfm

```

Our implementation of the monitor synthesizer is now complete. To use it, one can ask Maude reduce terms of the form `genMonitor(F)`, where `F` is the formula that one wants to generate into a monitor. For example:

```

reduce genMonitor(
  !('a /\ !(0 'b /\ 'c S ('d /\ (! 'e S 'f))))
) .

```

For the formula above, Maude will give the expected answer, pretty printed as follows:

```

\|||||/
--- Welcome to Maude ---
/|||||/
Maude 2.2 built: Mar 15 2006 16:37:22
Copyright 1997-2005 SRI International
Sat Jan 27 12:01:20 2007
Maude> in p
=====
fmod K
=====
fmod PREDICATE
=====
fmod SYNTAX
=====
fmod CODE
=====
fmod MONITOR-GENERATION
=====
fmod PRETTY-PRINT

```

```

=====
reduce in PRETTY-PRINT :
  genMonitor(! ('a /\ ! (0 'b /\ 'c S ('d /\ ! 'e S 'f)))) .
rewrites: 46 in -93406740ms cpu (1ms real) (~ rewrites/second)
result Code:
bit[1 .. 3] := false ;
foreach new state s do {
  // "first update the bits in a consistent order"
  bit'[1] := 'b(s) ;
  bit'[2] := 'f(s) \/ ! 'e(s) /\ bit[2] ;
  bit'[3] := 'd(s) /\ bit'[2] \/ 'c(s) /\ bit[3] ;
  // "then check whether the formula is violated"
  if 'a(s) /\ ! (bit[1] /\ bit'[3]) then
    Error ;
  // "finally, update the state of the monitor"
  bit[1 .. 3] := bit'[1 .. 3]
}

Maude>

```

## The K Module

One should upload the next module whenever one wants to use the K technique to define a language, logic or tool. Note that the module below has nothing to do with our particular logic under consideration in this paper; that is the reason for which we exiled it here.

```

fmod K is
  sorts KConfigurationItem KConfiguration .
  subsort KConfigurationItem < KConfiguration .
  op empty : -> KConfiguration .
  op __ : KConfiguration KConfiguration -> KConfiguration [assoc comm id: empty] .

  sorts KComputationItem KNeComputation KComputation .
  subsort KComputationItem < KNeComputation < KComputation .
  op nil : -> KComputation .
  op _->_ : KComputation KComputation -> KComputation [assoc id: nil] .
  op _->_ : KNeComputation KNeComputation -> KNeComputation [ditto] .

  sort KComputationList .
  subsort KComputation < KComputationList .
  op nil : -> KComputationList .
  op _,_ : KComputationList KComputationList -> KComputationList [assoc id: nil] .

  sort KResult KResultList .
  subsorts KResult < KResultList < KComputation .

```

```

op nil : -> KResultList .
op _,_ : KResultList KResultList -> KResultList [assoc id: nil] .

op [_] : KComputationList -> KComputationItem .
op [_] : KResultList -> KComputationItem .
op [_|_] : KComputationList KResultList -> KComputationItem .

var K : KNeComputation . var Kl : KComputationList .
var R : KResult . var Rl : KResultList .
eq [K,Kl] = K -> [Kl | nil] .
eq R -> [K,Kl | Rl] = K -> [Kl | Rl,R] .
eq R -> [nil | Rl] = [Rl,R] .
endfm

```

To use K, after importing the module above, one should define one's own constructors for configuration items (sort `KConfigurationItem`), for computation items (sort `KComputationItem`), and for results (sort `KResult`). For our example, we defined all these at the beginning of the module `MONITOR-GENERATION`.



## Chapter 11

# Parametric Property Monitoring



## Chapter 12

# Predictive Runtime Analysis



## Chapter 13

# Static Analysis to Improve Runtime Verification



## Chapter 14

# Semantics-Based Runtime Verification

14.1 Defining a Formal Semantics

14.2 Semantics-Based Symbolic Execution

14.3 Program Verification as Exhaustive Runtime Verification



## Chapter 15

# Conclusion and Future Work

### 15.1 Safety Properties and Monitoring

Chapters 3 and 4 presented a comprehensive study of safety properties and of their monitoring, using a uniform formalism and notation. Technically, there were two novel contributions. First, it introduced the notion of a *persistent* safety property, which is the finite-trace correspondent of an infinite-trace safety property, and used it to show the cardinal equivalence of the various notions of safety property encountered in the literature. Second, it rigorously defined the problem of monitoring a safety property, and it showed that it can be arbitrarily hard. These results established a firm foundation for studying safety properties and corresponding monitors and algorithms for various domains of interest, where requirements can be expressed using domain-specific formalisms, such as future-time and past-time temporal logics, context-free grammars, push-down automata, and so on.



# Bibliography

- [1] MOP website, LTL plugin. <http://fsl.cs.uiuc.edu/mop/logic-plugins/ltl>.
- [2] Martín Abadi and Leslie Lamport. The existence of refinement mappings. In *LICS*, pages 165–175. IEEE Computer Society, 1988.
- [3] Martín Abadi and Leslie Lamport. The existence of refinement mappings. *Theoretical Computer Science*, 82(2):253–284, 1991.
- [4] Chris Allan, Pavel Avgustinov, Aske Simon Christensen, Laurie Hendren, Sascha Kuzins, Ondrej Lhoták, Oege de Moor, Damien Sereni, Ganesh Sittampalam, and Julian Tibble. Adding trace matching with free variables to AspectJ. In Richard P. Gabriel, editor, *ACM Conference on Object-Oriented Programming, Systems and Languages (OOPSLA)*, pages 345–364. ACM Press, 2005.
- [5] Bowen Alpern and Fred B. Schneider. Defining liveness. *IPL*, 21(4):181–185, 1985.
- [6] V. M. Antimirov. Partial derivatives of regular expressions and finite automaton constructions. *Theoretical Computer Science*, 155(2):291–319, 1996.
- [7] V. M. Antimirov and P. D. Mosses. Rewriting extended regular expressions. *Theoretical Computer Science*, 143(1):51–72, 1995.
- [8] Krzysztof R. Apt, Nissim Francez, and Shmuel Katz. Appraising fairness in languages for distributed programming. In *POPL*, pages 189–198, 1987.
- [9] Cyrille Artho, Doron Drusinsky, Allen Goldberg, Klaus Havelund, Mike Lowry, Corina Pasareanu, Grigore Roşu, and Willem Visser. Experiments with Test Case Generation and Runtime Analysis. In *Proc. of ASM’03: Abstract State Machines*, volume 2589 of *Lecture Notes in Computer Science*, pages 87–107, Taormina, Italy, March 2003. Springer.
- [10] Cyrille Artho, Doron Drusinsky, Allen Goldberg, Klaus Havelund, Mike Lowry, Corina Păsăreanu, Grigore Roşu, Willem Visser, and Rich Washington.

Automated Testing using Symbolic Execution and Temporal Monitoring. *Theoretical Computer Sci.*, to appear, 2005.

- [11] T. Ball, A. Podelski, and S. Rajamani. Boolean and Cartesian Abstractions for Model Checking C Programs. In *Proc. of TACAS'01: Tools and Algorithms for the Construction and Analysis of Systems*, Lecture Notes in Computer Science, Genova, Italy, April 2001.
- [12] H. Barringer, A. Goldberg, K. Havelund, and K. Sen. Rule-Based Runtime Verification. In *Proceedings of VMCAI'04*, volume 2937 of *LNCS*, pages 44–57, 2004.
- [13] Adel Bouhoula, Jean-Pierre Jouannaud, and José Meseguer. Specification and Proof in Membership Equational Logic. *Theoretical Computer Science*, 236:35–132, 2000.
- [14] R. E. Bryant. Graph-Based Algorithms for Boolean Function Manipulation. *IEEE Transactions on Computers*, C-35(8):677–691, August 1986.
- [15] J R Büchi. *On a Decision Method in Restricted Second Order Arithmetic*. Logic, Methodology and Philosophy of Sciences. Stanford University Press, 1962.
- [16] A. K. Chandra, D. C. Kozen, and L. J. Stockmeyer. Alternation. *Journal of the ACM*, 28(1):114–133, 1981.
- [17] Ashok K. Chandra, Dexter Kozen, and Larry J. Stockmeyer. Alternation. *J. ACM*, 28(1):114–133, 1981.
- [18] F. Chen and G. Roşu. Towards Monitoring-Oriented Programming: A Paradigm Combining Specification and Implementation. In *Proc. of RV'03: the Third International Workshop on Runtime Verification*, volume 89 of *Electronic Notes in Theoretical Computer Science*, pages 106–125, Boulder, Colorado, USA, 2003. Elsevier Science.
- [19] Feng Chen, Marcelo d'Amorim, and Grigore Roşu. A Formal Monitoring-Based Framework for Software Development and Analysis. In *Proceedings of ICFEM'04*, volume 3308 of *LNCS*, pages 357–372, 2004.
- [20] Feng Chen, Marcelo D'Amorim, and Grigore Roşu. Checking and correcting behaviors of Java programs at runtime with Java-MOP. In *RV'05*, volume 144(4) of *ENTCS*, 2005.
- [21] Edmund M. Clarke, Orna Grumberg, and Doron A. Peled. *Model Checking*. The MIT Press, Cambridge, Massachusetts, 1999.

- [22] M. Clavel. The ITP Tool. In *Logic, Language and Information. Proc. of the First Workshop on Logic and Language*, pages 55–62. Kronos, 2001.
- [23] M. Clavel, F. Durán, S. Eker, P. Lincoln, N. Martí-Oliet, J. Meseguer, and J. F. Quesada. Towards Maude 2.0. In *3rd International Workshop on Rewriting Logic and its Applications (WRLA'00)*, volume 36 of *Electronic Notes in Theoretical Computer Science*. Elsevier, 2000.
- [24] Manuel Clavel, Francisco Durán, Steven Eker, Patrick Lincoln, Narciso Martí-Oliet, José Meseguer, and Jose F. Quesada. Maude: specification and programming in rewriting logic. *Theoretical Computer Science*, 285(2):187–243, 2002.
- [25] Manuel Clavel, Francisco Durán, Steven Eker, Patrick Lincoln, Narciso Martí-Oliet, José Meseguer, and Carolyn Talcott. Maude Manual. <http://maude.cs.uiuc.edu>.
- [26] Manuel Clavel, Francisco J. Durán, Steven Eker, Patrick Lincoln, Narciso Martí-Oliet, José Meseguer, and José F. Quesada. Maude: Specification and Programming in Rewriting Logic. *Theoretical Computer Science*, 285:187–243, 2002.
- [27] James Corbett, Matthew B. Dwyer, John Hatcliff, Corina S. Pasareanu, Robby, Shawn Laubach, and Hongjun Zheng. Bandera : Extracting Finite-state Models from Java Source Code. In *Proc. of ICSE'00: International Conference on Software Engineering*, Limerich, Ireland, June 2000. ACM Press.
- [28] M. Dahm. BCEL. <http://jakarta.apache.org/bcel>.
- [29] Marcelo d'Amorim and Grigore Roşu. Efficient monitoring of  $\omega$ -languages. In Kousha Etessami and Sriram K. Rajamani, editors, *CAV*, volume 3576 of *LNCS*, pages 364–378, 2005.
- [30] Marcelo d'Amorim and Grigore Roşu. Efficient monitoring of  $\omega$ -languages. In *Proceedings of 17th International Conference on Computer-aided Verification (CAV'05)*, volume 3576 of *Lecture Notes in Computer Science*, pages 364 – 378. Springer, 2005.
- [31] Claudio Demartini, Radu Iosif, and Riccardo Sisto. A Deadlock Detection Tool for Concurrent Java Programs. *Software Practice and Experience*, 29(7):577–603, July 1999.
- [32] D. Drusinsky. The Temporal Rover and the ATG Rover. In *SPIN Model Checking and Software Verification*, volume 1885 of *Lecture Notes in Computer Science*, pages 323–330. Springer, 2000.

- [33] Doron Drusinsky. Monitoring Temporal Rules Combined with Time Series. In *Proc. of CAV'03: Computer Aided Verification*, volume 2725 of *Lecture Notes in Computer Science*, pages 114–118, Boulder, Colorado, USA, 2003. Springer-Verlag.
- [34] Kousha Etessami and Gerard Holzmann. Optimizing Büchi Automata. In *Proc. of Int. Conf. on Concurrency Theory*, volume 1877 of *LNCS*, pages 153–167. Springer, 2000.
- [35] James Ezick. An Optimizing Compiler for Batches of Temporal Logic Formulas. In *Proceedings of ISSTA'04*, pages 183–194. ACM Press, 2004.
- [36] C. J. Fidge. Partial Orders for Parallel Debugging. In *Proc. of the 1988 ACM SIGPLAN and SIGOPS workshop on Parallel and Distributed Debugging*, pages 183–194. ACM, 1988.
- [37] B. Finkbeiner and H. Sipma. Checking Finite Traces using Alternating Automata. In *Proc. of RV'01: The First International Workshop on Runtime Verification*, volume 55(2) of *Electronic Notes in Theoretical Computer Science*, Paris, France, 2001. Elsevier Science.
- [38] Bernd Finkbeiner, Sriram Sankaranarayanan, and Henny Sipma. Collecting Statistics over Runtime Executions. In *Proc. of RV'02: The Second International Workshop on Runtime Verification*, volume 70 of *Electronic Notes in Theoretical Computer Science*, Paris, France, 2002. Elsevier.
- [39] E. R. Gansner and S. C. North. An open graph visualization system and its applications to software engineering. *Software Practice and Experience*, 30(1):1203–1233, September 2000.
- [40] Michael Garey. Optimal Binary Identification Procedures. *SIAM Journal on Applied Mathematics*, 23(2):173–186, 1972.
- [41] P. Gastin and D. Oddoux. LTL with Past and Two-Way Very-Weak Alternating Automata. In *Proceedings of MFCS'03*, number 2747 in *LNCS*, pages 439–448. Springer, 2003.
- [42] Marc Geilen. On the Construction of Monitors for Temporal Logic Properties. In *Proceedings of RV'01*, volume 55 of *ENTCS*. Elsevier Science, 2001.
- [43] Rob Gerth, Doron Peled, Moshe Vardi, and Pierre Wolper. Simple On-the-fly Automatic Verification of Linear Temporal Logic. In *Proc. of the 15th Workshop on Protocol Specification, Testing, and Verification*. North-Holland, 1995.

- [44] D. Giannakopoulou and K. Havelund. Automata-Based Verification of Temporal Properties on Running Programs. In *Proceedings, International Conference on Automated Software Engineering (ASE'01)*, pages 412–416. Institute of Electrical and Electronics Engineers, 2001. Coronado Island, California.
- [45] Patrice Godefroid. Model Checking for Programming Languages using VeriSoft. In *Proc. of POPL'97: the 24th ACM Symposium on Principles of Programming Languages*, pages 174–186, Paris, France, January 1997.
- [46] J. Goguen, K. Lin, and G. Roşu. Circular coinductive rewriting. In *Proceedings, Automated Software Engineering '00*, pages 123–131. IEEE, 2000. (Grenoble, France).
- [47] J. Goguen, K. Lin, and G. Rosu. Conditional circular coinductive rewriting with case analysis. In *Recent Trends in Algebraic Development Techniques (WADT'02)*, Lecture Notes in Computer Science, to appear, Frauenchiemsee, Germany, September 2002. Springer-Verlag.
- [48] Joseph Goguen, Kai Lin, Grigore Roşu, Akira Mori, and Bogdan Warinschi. An overview of the Tatami project. In *Cafe: An Industrial-Strength Algebraic Formal Method*, pages 61–78. Elsevier, 2000.
- [49] Joseph Goguen, James Thatcher, Eric Wagner, and Jesse Wright. Initial Algebra Semantics and Continuous Algebras. *Journal of the Association for Computing Machinery*, 24(1):68–95, January 1977.
- [50] Joseph Goguen, Timothy Winkler, José Meseguer, Kokichi Futatsugi, and Jean-Pierre Jouannaud. Introducing OBJ. In *Software Engineering with OBJ: Algebraic Specification in Action*. Kluwer, 2000.
- [51] Elsa Gunter and Doron Peled. Tracing the Executions of Concurrent Programs. In *Proc. of RV'02: Second International Workshop on Runtime Verification*, volume 70 of *Electronic Notes in Theoretical Computer Science*, Copenhagen, Denmark, 2002. Elsevier.
- [52] Elsa L. Gunter, Robert P. Kurshan, and Doron Peled. PET: An Interactive Software Testing Tool. In *Proc. of CAV'00: Computer Aided Verification*, volume 1885 of *Lecture Notes in Computer Science*, pages 552–556, Chicago, Illinois, USA, 2003. Springer-Verlag.
- [53] Elsa L. Gunter and Doron Peled. Using Functional Languages in Formal Methods: The PET System. In *Parallel and Distributed Processing Techniques and Applications*, pages 2981–2986. CSREA, 2000.
- [54] Kevin W. Hamlen, J. Gregory Morrisett, and Fred B. Schneider. Computability classes for enforcement mechanisms. *ACM Trans. Program. Lang. Syst.*, 28(1):175–205, 2006.

- [55] K. Havelund and G. Roşu. Java PathExplorer – A Runtime Verification Tool. In *The 6th International Symposium on Artificial Intelligence, Robotics and Automation in Space: A New Space Odyssey*, Montreal, Canada, June 18 - 21, 2001.
- [56] K. Havelund and G. Roşu. Monitoring Java Programs with Java PathExplorer. In *Proceedings of Runtime Verification (RV'01)*, volume 55 of *Electronic Notes in Theoretical Computer Science*. Elsevier Science, 2001.
- [57] K. Havelund and G. Roşu. Monitoring Programs using Rewriting. In *Proceedings, International Conference on Automated Software Engineering (ASE'01)*, pages 135–143. Institute of Electrical and Electronics Engineers, 2001. Coronado Island, California.
- [58] K. Havelund and G. Roşu. *Runtime Verification 2002*, volume 70(4) of *Electronic Notes in Theoretical Computer Science*. Elsevier Science, 2002. Proceedings of a *Computer Aided Verification (CAV'02)* satellite workshop.
- [59] K. Havelund and G. Roşu. Efficient monitoring of safety properties. *Software Tools and Technology Transfer*, 6(2):158–173, 2004. (also TACAS'02, LNCS 2280).
- [60] K. Havelund and G. Roşu. Synthesizing monitors for safety properties. In *Tools and Algorithms for Construction and Analysis of Systems (TACAS'02)*, volume 2280 of *Lecture Notes in Computer Science*, pages 342–356. Springer, 2002.
- [61] Klaus Havelund, Scott Johnson, and Grigore Roşu. Specification and Error Pattern Based Program Monitoring. In *Proc. of the European Space Agency workshop on On-Board Autonomy*, Noordwijk, The Netherlands, October 2001.
- [62] Klaus Havelund and Thomas Pressburger. Model Checking Java Programs using Java PathFinder. *International Journal on Software Tools for Technology Transfer*, 2(4):366–381, April 2000. Special issue containing selected submissions to SPIN'98: the fourth SPIN workshop, Paris, France, 1998.
- [63] Klaus Havelund and Grigore Roşu. Java PathExplorer – A Runtime Verification Tool. In *Proc. of i-SAIRAS'01: the 6th International Symposium on Artificial Intelligence, Robotics and Automation in Space*, Montreal, Canada, June 2001.
- [64] Klaus Havelund and Grigore Roşu. Monitoring Java Programs with Java PathExplorer. In *Proc. of RV'01: the First International Workshop on Runtime Verification*, volume 55 of *Electronic Notes in Theoretical Computer Science*, pages 97–114, Paris, France, July 2001. Elsevier Science.

- [65] Klaus Havelund and Grigore Roşu. *Workshops on Runtime Verification (RV'01, RV'02, RV'04)*, volume 55, 70(4), to appear of *ENTCS*. Elsevier, 2001, 2002, 2004.
- [66] S. Hirst. A new algorithm solving membership of extended regular expressions. Technical report, The University of Sydney, 1989.
- [67] Gerard Holzmann. The Model Checker SPIN. *IEEE Transactions on Software Engineering*, 23(5):279–295, 1997.
- [68] Gerard J. Holzmann and Margaret H. Smith. A Practical Method for Verifying Event-Driven Software. In *Proc. of ICSE'99: International Conference on Software Engineering*, Los Angeles, California, USA, May 1999. IEEE/ACM.
- [69] J. E. Hopcroft and J. D. Ullman. *Introduction to Automata Theory, Languages and Computation*. Addison Wesley, 1979.
- [70] J.E. Hopcroft and J.D. Ullman. *Introduction to Automata Theory, Languages, and Computation*. Addison-Wesley, 1979.
- [71] Jieh Hsiang. Refutational Theorem Proving using Term Rewriting Systems. *Artificial Intelligence*, 25:255–300, 1985.
- [72] Laurent Hyafil and Ronald L. Rivest. Computing Optimal Binary Decision Trees is NP-complete. *Information Processing Letters*, 5(1):15–17, 1976.
- [73] L. Ilie, B. Shan, and S. Yu. Fast algorithms for extended regular expression matching and searching. In *Proceedings of STACS'03*, volume 2607 of *LNCS*, pages 179–190, 2003.
- [74] L. Ilie, B. Shan, and S. Yu. Fast algorithms for extended regular expression matching and searching. In H. Alt and M. Habib, editors, *Proceedings of the 20th International Symposium on Theoretical Aspects of Computer (STACS 03)*, volume 2607 of *Lecture Notes in Computer Science*, page 179. Springer-Verlag, Berlin, 2003.
- [75] Yonit Kesten, Zohar Manna, Hugh McGuire, and Amir Pnueli. A Decision Algorithm for Full Propositional Temporal Logic. In *Proceedings of CAV'93*, volume 697 of *LNCS*, pages 97–109. Springer, 1993.
- [76] M. Kim, S. Kannan, I. Lee, and O. Sokolsky. Java-MaC: a Run-time Assurance Tool for Java. In *Proceedings of Runtime Verification (RV'01)*, volume 55 of *Electronic Notes in Theoretical Computer Science*. Elsevier Science, 2001.
- [77] J.R. Knight and E.W. Myers. Super-pattern matching. *Algorithmica*, 13(1/2):211–243, 1995.

- [78] David Kortenkamp, Tod Milam, Reid Simmons, and Joaquin Fernandez. Collecting and Analyzing Data from Distributed Control Programs. In *Proc. of RV'01: First International Workshop on Runtime Verification*, volume 55 of *Electronic Notes in Theoretical Computer Science*, Paris, France, 2001. Elsevier Science.
- [79] O. Kupferman and M. Y. Vardi. Freedom, Weakness, and Determinism: From Linear-Time to Branching-Time. In *Proc. of the IEEE Symposium on Logic in Computer Science*, pages 81–92, 1998.
- [80] O. Kupferman and M. Y. Vardi. Model Checking of Safety Properties. In *Proc. of CAV'99: Conference on Computer-Aided Verification*, Trento, Italy, 1999.
- [81] O. Kupferman and S. Zuhovitzky. An Improved Algorithm for the Membership Problem for Extended Regular Expressions. In *Proc. of the International Symposium on Mathematical Foundations of Computer Science*, volume 2420 of *Lecture Notes in Computer Science*, 2002.
- [82] O. Kupferman and S. Zuhovitzky. An improved algorithm for the membership problem for extended regular expressions. In *Proc. of MFCS'02*, volume 2420 of *LNCS*, pages 446–458, 2002.
- [83] Orna Kupferman and Moshe Y. Vardi. Model checking of safety properties. *Formal Methods in System Design*, 19(3):291–314, 2001.
- [84] Orna Kupferman and Moshe Y. Vardi. Model checking of safety properties. *Formal Methods in System Design*, 19(3):291–314, October 2001.
- [85] Leslie Lamport. Proving the correctness of multiprocess programs. *IEEE Trans. Software Eng.*, 3(2):125–143, 1977.
- [86] Leslie Lamport. Logical foundation. In M. W. Alford, J. P. Ansart, G. Hommel, L. Lamport, B. Liskov, G. P. Mullery, F. B. Schneider, M. Paul, and H. J. Siebert, editors, *Distributed systems: Methods and tools for specification. An advanced course*, volume 190 of *LNCS*, pages 119–130. Springer-Verlag, 1985.
- [87] I. Lee, S. Kannan, M. Kim, O. Sokolsky, and M. Viswanathan. Runtime Assurance Based on Formal Specifications. In *Proceedings of the International Conference on Parallel and Distributed Processing Techniques and Applications*, 1999.
- [88] Zohar Manna and Amir Pnueli. *The Temporal Logic of Reactive and Concurrent Systems*. Springer, New York, 1992.

- [89] Zohar Manna and Amir Pnueli. *Temporal verification of reactive systems: safety*. Springer-Verlag New York, Inc., New York, NY, USA, 1995.
- [90] Zohar Manna and Amir Pnueli. *Temporal Verification of Reactive Systems: Safety*. Springer, New York, 1995.
- [91] Nicolas Markey. Temporal Logic with Past is Exponentially more Succinct. *EATCS Bulletin*, 79:122–128, 2003.
- [92] Nicolas Markey and Philippe Schnoebelen. Model Checking a Path (Preliminary Report). In *Proc. of CONCUR’03: International Conference on Concurrency Theory*, volume 2761 of *Lecture Notes in Computer Science*, pages 251–265, Marseille, France, August 2003. Springer.
- [93] José Meseguer. Conditional Rewriting Logic as a Unified Model of Concurrency. *Theoretical Computer Science*, pages 73–155, 1992.
- [94] José Meseguer. Membership Algebra as a Logical Framework for Equational Specification. In *Proc. of WADT’97: Workshop on Algebraic Development Techniques*, volume 1376 of *Lecture Notes in Computer Science*, pages 18–61, Tarquinia, Italy, June 1998. Springer.
- [95] Bernard Moret. Decision Trees and Diagrams. *ACM Comp. Surv.*, 14(4):593–623, 1982.
- [96] G. Myers. A four russians algorithm for regular expression pattern matching. *Journal of the ACM*, 39(4):430–448, 1992.
- [97] Dennis Oddoux and Paul Gastin. LTL2BA. <http://www.liafa.jussieu.fr/~oddoux/ltl2ba/>.
- [98] T. O’Malley, D. Richardson, and L. Dillon. Efficient Specification-Based Oracles for Critical Systems. In *In Proceedings of the California Software Symposium*, 1996.
- [99] David Y.W. Park, Ulrich Stern, and David L. Dill. Java Model Checking. In *Proc. of the First International Workshop on Automated Program Analysis, Testing and Verification*, Limerick, Ireland, June 2000.
- [100] Amir Pnueli. The Temporal Logic of Programs. In *Proc. of the 18th IEEE Symposium on Foundations of Computer Science*, pages 46–77, 1977.
- [101] D. J. Richardson, S. L. Aha, and T. O. O’Malley. Specification-Based Test Oracles for Reactive Systems. In *Proceedings of the Fourteenth International Conference on Software Engineering, Melbourne, Australia*, pages 105–118, 1992.

- [102] G. Roşu. *Hidden Logic*. PhD thesis, University of California at San Diego, 2000.
- [103] G. Roşu and M. Viswanathan. Testing extended regular language membership incrementally by rewriting. In *RTA'03*, volume 2706 of *LNCS*. Springer, 2003.
- [104] Grigore Roşu and Klaus Havelund. Synthesizing Dynamic Programming Algorithms from Linear Temporal Logic Formulae. RIACS Technical report TR 01-08, January 2001.
- [105] Grigore Roşu and Klaus Havelund. Rewriting-based techniques for runtime verification. *Automated Software Engineering*, 12(2):151–197, 2005.
- [106] Grigore Roşu and Klaus Havelund. Rewriting-Based Techniques for Runtime Verification. *Journal of Automated Software Engineering*, 2005. to appear.
- [107] Grigore Roşu. K: a rewrite-based framework for modular lang. design, semantics, analysis and implementation (V2). Technical Report UIUCDCS-R-2006-2802, 2006.
- [108] Grigore Rosu. An effective algorithm for the membership problem for extended regular expressions. In *Proceedings of the 10th International Conference on Foundations of Software Science and Computation Structures (FOSSACS'07)*, volume 4423 of *LNCS*, pages 332–345. Springer-Verlag, 2007.
- [109] Grigore Rosu. On safety properties and their monitoring. *Sci. Ann. Comp. Sci.*, 22(2):327–365, 2012.
- [110] J. J. M. M. Rutten. Automata and coinduction (an exercise in coalgebra). In *Proceedings of the 9th International Conference on Concurrency Theory (CONCUR 98)*, volume 1466 of *Lecture Notes in Computer Science*, pages 194–218. Springer-Verlag, 1998.
- [111] Stefan Savage, Michael Burrows, Greg Nelson, Patrik Sobalvarro, and Thomas Anderson. Eraser: A Dynamic Data Race Detector for Multithreaded Programs. *ACM Transactions on Computer Systems*, 15(4):391–411, November 1997.
- [112] Fred B. Schneider. *On Concurrent Programming*. Springer, 1997.
- [113] Fred B. Schneider. Enforceable security policies. *ACM Trans. Inf. Syst. Secur.*, 3(1):30–50, 2000.
- [114] Alper Sen and Vijay K. Garg. Partial Order Trace Analyzer (POTA) for Distributed Programs. In *Proc. of RV'03: the Third International Workshop on Runtime Verification*, volume 89 of *Electronic Notes in Theoretical Computer Science*, Boulder, Colorado, USA, 2003. Elsevier Science.

- [115] K. Sen and G. Roşu. Generating Optimal Monitors for Extended Regular Expressions. In *Proc. of RV'03: the Third International Workshop on Runtime Verification*, volume 89 of *Electronic Notes in Theoretical Computer Science*, pages 162–181, Boulder, Colorado, USA, 2003. Elsevier Science.
- [116] K. Sen, G. Roşu, and G. Agha. Runtime safety analysis of multithreaded programs. Technical Report UIUCDCS-R-2003-2334, University of Illinois at Urbana Champaign, April 2003.
- [117] Koushik Sen, Grigore Roşu, and Gul Agha. Runtime Safety Analysis of Multithreaded Programs. In *Proc. of ESEC/FSE'03: European Software Engineering Conference and ACM SIGSOFT International Symposium on the Foundations of Software Engineering*. ACM, Helsinki, Finland, September 2003.
- [118] Koushik Sen and Grigore Rosu. Generating optimal monitors for extended regular expressions. In *Proceedings of 3rd International Workshop on Runtime Verification (RV'03)*, volume 89(2) of *Electronic Notes in Theoretical Computer Science*, pages 226–245. Elsevier, June 2003.
- [119] Koushik Sen, Grigore Roşu, and Gul Agha. Online Efficient Predictive Safety Analysis of Multithreaded Programs. In *Proceedings of TACAS'04*, volume 2988 of *LNCS*, pages 123–138. Springer, 2002.
- [120] Natarajan Shankar, Sam Owre, and John M. Rushby. *PVS Tutorial*. Computer Science Laboratory, SRI International, Menlo Park, CA, February 1993. Also appears in Tutorial Notes, *Formal Methods Europe '93: Industrial-Strength Formal Methods*, pages 357–406, Odense, Denmark, April 1993.
- [121] A. P. Sistla and E. M. Clarke. The Complexity of Propositional Linear Temporal Logics. *Journal of the ACM (JACM)*, 32(3):733–749, 1985.
- [122] A. P. Sistla and E. M. Clarke. The Complexity of Propositional Linear Temporal Logics. *Journal of the ACM*, 32(3):733–749, 1985.
- [123] O. Sokolsky and M. Viswanathan. *Workshop on Runtime Verification (RV'03)*, volume 89 of *ENTCS*. Elsevier, 2003.
- [124] L. J. Stockmeyer and A. R. Meyer. Word problems requiring exponential time (preliminary report). In *STOC*, pages 1–9. ACM Press, 1973.
- [125] Larry Joseph Stockmeyer. *The Complexity of Decision Problems in Automata Theory and Logic*. PhD thesis, Massachusetts Institute of Technology, 1974.
- [126] Scott D. Stoller. Model-Checking Multi-Threaded Distributed Java Programs. In *Proc. of SPIN'00: SPIN Model Checking and Software Verification*, volume 1885 of *Lecture Notes in Computer Science*, pages 224–244, Stanford, California, USA, 2000. Springer.

- [127] K. Thompson. Regular expression search algorithm. *CACM*, 11(6):419–422, 1968.
- [128] Willem Visser, Klaus Havelund, Guillaume Brat, and SeungJoon Park. Model Checking Programs. In *Proc. of ASE'00: International Conference on Automated Software Engineering*, Grenoble, France, September 2000. IEEE CS Press.
- [129] Pierre Wolper. Constructing Automata from Temporal Logic Formulas: a Tutorial. volume 2090 of *LNCS*, pages 261–277. Springer, 2002.
- [130] H. Yamamoto. An automata-based recognition algorithm for semi-extended regular expressions. In *Proc. of MFCS'00*, volume 1893 of *LNCS*, pages 699–708, 2000.
- [131] H. Yamamoto. A new recognition algorithm for extended regular expressions. In *Proceedings of ISAAC'01*, volume 2223 of *LNCS*, pages 257–267, 2001.
- [132] H. Yamamoto and T. Miyazaki. A fast bit-parallel algorithm for matching extended regular expressions. In *Proc. of COCOON'03*, volume 2697 of *LNCS*, pages 222–231, 2003.